

BrainXPlore

Decision finding in Brain Biopsy Planning

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Visual Computing

eingereicht von

Lukas Pezenka

Matrikelnummer 0625948

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Univ.-Doz. Dipl.-Ing. Dr.techn. Eduard Gröller

Mitwirkung: Dipl.-Math. Dr. Katja Bühler

Assoc.Prof. Dr. Stefan Wolfsberger

Wien, 28. November 2017

Lukas Pezenka

Eduard Gröller

BrainXPlore

Decision finding in Brain Biopsy Planning

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Visual Computing

by

Lukas Pezenka

Registration Number 0625948

to the Faculty of Informatics

at the TU Wien

Advisor: Ao.Univ.Prof. Univ.-Doz. Dipl.-Ing. Dr.techn. Eduard Gröller

Assistance: Dipl.-Math. Dr. Katja Bühler

Assoc.Prof. Dr. Stefan Wolfsberger

Vienna, 28th November, 2017

Lukas Pezenka

Eduard Gröller

Erklärung zur Verfassung der Arbeit

Lukas Pezenka
Gerasdorfer Straße 61/2/6

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 28. November 2017

Lukas Pezenka

Acknowledgements

This work was enabled by the Competence Centre VRVis. VRVis is funded by BMVIT, BMWF, Styria, SFG and Vienna Business Agency in the scope of COMET - Competence Centers for Excellent Technologies (854174) which is managed by FFG.

I thank my partner, Sabine, for her patience and support. Further, I thank my parents for encouraging me to pursue an academic education. Finally, thank my friends and colleagues, especially Chris and Nicolas, for providing feedback and offering advice whenever I was stuck.

A big salute goes to my supervisors, Katja Bühler and Eduard Gröller for supervision and proofreading of this thesis. Additional thanks go to Florian, Alexey and Gert for their advice and support.

Kurzfassung

Neurochirurgen treffen Entscheidungen auf Basis ihres Expertenwissens. In diesem Prozess werden Faktoren wie die Distanz zu gefährdeten Strukturen, Pfadlänge und -winkel, etc. berücksichtigt. Während manche Kriterien (wie etwa das Unterschreiten der minimalen Distanz zu Blutgefäßen) zum Ausschluß einer potentiellen Trajektorie führen, sind andere Kriterien (etwa die Pfadlänge) weniger rigide. Diese Faktoren können dann zur Reihung und Bewertung von potentiellen Trajektorien herangezogen werden. Nach Analyse tatsächlich erfolgter Eingriffe sowie Rücksprache mit Neurochirurgen haben wir wichtige Regeln zur Planung und Durchführung von Gehirnbiopsien identifiziert.

In dieser Arbeit präsentieren wir BRAINXPLORE, ein in Zusammenarbeit mit der medizinischen Universität Wien entwickeltes System zur Unterstützung von Neurochirurgen bei der Planung von Gehirnbiopsien. BRAINXPLORE ist ein erweiterbares Biopsieplanungsrahmenwerk, das die erarbeiteten Regeln implementiert und dabei dem Benutzer volle Flexibilität hinsichtlich deren Definition und Erweiterung bietet. Das System berechnet automatisch eine Menge an potentiellen Trajektorien. Durch die Definition und Verfeinerung der Regeln kann diese Menge schrittweise verkleinert werden, bis eine geeignete Trajektorie gefunden wurde. Durch den Einsatz eines räumlichen Indexes können beliebig viele anatomische Strukturen bei beliebiger Auflösung zur Berechnung der Trajektorien herangezogen werden. Somit wird der Einsatz hochauflösender multimodaler Datensätze möglich, welcher bisher aufgrund von Speicherplatzlimitierungen der Graphikkarte unmöglich war. Um den Neurochirurgen bei der Entscheidungsfindung zu unterstützen bietet BrainXPlore Methoden der Informationsvisualisierung, wie parallele Koordinatensysteme und Risikoprofilogramme.

Wir haben BRAINXPLORE auf realen Daten einer tatsächlich erfolgten Biopsie getestet und akzeptable Resultate erzielt. Der partizipierende Neurochirurg gab uns das Feedback, dass der Einsatz von BRAINXPLORE zu einer Verringerung der Dauer von Gehirnbiopsien führen kann und sinnvoll zur Unterstützung weniger erfahrener Chirurgen eingesetzt werden kann.

Abstract

Neurosurgeons make decisions based on expert knowledge that takes factors such as safety margins, the avoidance of risk structures, trajectory length and trajectory angle into consideration. While some of those factors are mandatory, others can be optimized in order to obtain the best possible trajectory under the given circumstances. Through comparison with the actually chosen trajectories from real biopsies and qualitative interviews with domain experts, we identified important rules for trajectory planning.

In this thesis, we present BRAINXPLORE, an interactive visual analysis tool for aiding neurosurgeons in planning brain biopsies. BRAINXPLORE is an extendable Biopsy Planning framework that incorporates those rules while at the same time leaving full flexibility for their customization and adding of new structures at risk. Automatically computed candidate trajectories can be incrementally refined in an interactive manner until an optimal trajectory is found. We employ a spatial index server as part of our system that allows us to access distance information on an unlimited number of risk structures at arbitrary resolution. Furthermore, we implemented InfoVis techniques such as Parallel Coordinates and risk signature charts to drive the decision process. As a case study, BRAINXPLORE offers a variety of information visualization modalities to present multivariate data in different ways.

We evaluated BRAINXPLORE on a real dataset and accomplished acceptable results. The participating neurosurgeon gave us the feedback that BRAINXPLORE can decrease the time needed for biopsy planning and aid novice users in their decision making process.

Contents

Kurzfassung	ix
Abstract	xi
Contents	xiii
1 Introduction	1
2 Medical Background	5
2.1 Medical Imaging Methods	5
2.2 Brain Anatomy Segmentation	12
2.3 Planning Guidelines	15
3 Related Work	19
3.1 Biopsy Planning	19
3.2 Medical Information Visualization	30
4 Visual Computing Work-flow and System Design	35
4.1 Requirement Analysis	35
4.2 System Overview	41
4.3 Visual Computing Work-flow	43
5 Methods and Data	45
5.1 Data Acquisition and Segmentation	45
5.2 Data Preprocessing	46
5.3 Indexing	46
5.4 Trajectory Generation and Preprocessing	47
5.5 Trajectory Filtering and Refinement	48
6 Implementation	55
6.1 MeVisLab Work-Flow	55
6.2 Data Preprocessing	56
6.3 Indexing	57
6.4 User Interface	57
	xiii

6.5	Information Visualization	67
7	Evaluation	71
7.1	Internal Evaluation	71
7.2	External Evaluation	86
8	Conclusion	91
8.1	Future Work	92
	List of Figures	95
	List of Tables	97
	List of Algorithms	99
	Bibliography	101

Introduction

More than 100 different tumor types with widely divergent biologies and clinical outcomes are known [SR06]. Different treatment options, such as surgery, radiation, and palliative cytotoxics exist and are well established in medicine. Before deciding between further treatment options, however, the tumor needs to be identified.

A brain tumor biopsy is the extraction of tissue from a tumor for the analysis of its nature, i.e., whether it is benign or malignant. For the biopsy, a hole is drilled into the cranium, through which a biopsy needle is inserted into the brain. In general, a needle trajectory is defined by an entry point in the skin and a target point on the tumor.

To prevent permanent damage to the patient, vital tissues have to be avoided by the trajectory and hence, protected from damage caused by the needle. Certain rules have to be fulfilled for a trajectory to be deemed acceptable. Amongst those rules are the avoidance of risk structures, minimization of trajectory length, etc.

Satisfying all those rules requires good knowledge of the particular brain's anatomy. Due to the high risk, carrying out the biopsy also requires a high degree of precision. A **stereotactic brain biopsy** makes use of medical imaging data such as magnetic resonance imaging (MRI) and computer tomography (CT) to precisely identify and target any brain area in a 3D coordinate system. The biopsy needle can be positioned by a robotic arm after the patient was registered to the 3D coordinate system. Stereotactic biopsies are well established and have a high success rate. Shoshan et al. [SPS⁺95] report a diagnostic success rate of 96% over a three-year period.

Deep Brain Stimulation (DBS) is a method of treating advanced Parkinson's disease [BSB04], movement disorders and other psychiatric and neurological disorders such as epilepsy [LKG14]. The neurosurgeon places electrodes in the patient's brain along well-defined neural pathways. These electrodes further serve to regulate the flow of neural information, including pathological signals mediating seizure propagation. This mech-

anism is known as neuro-modulation. Like brain biopsies, the placement of electrodes requires careful and precise planning.

Despite the available technologies and tools, brain biopsy and DBS planning is still a time-intensive task that requires a high amount of expert knowledge. Different systems have been proposed to plan and visualize needle trajectories. While some aspects of the planning process are (partially) solved by the proposed systems, others need to be improved upon.

For one, most systems operate on a rigid set of rules to augment the planning process. To our knowledge, no existing work allows the inclusion of an arbitrary number of distinct structures (e.g., fiber bundles, anatomical atlas, etc) and the flexible definition of rules for trajectory filtering. This does not necessarily allow neurosurgeons, who ultimately are the domain experts, to use the tool in such a way as to optimally support their (or their hospital's) preferred planning methods.

On the other hand, only a limited number of structures is taken into consideration by existing solutions at this moment. As brain atlases become more sophisticated and white matter tracts are better understood, the number of structures that are considered when planning a biopsy is likely to increase greatly in the near future. Orringer et al. [OGJ12] points out that advances in MRI imaging provide the possibility of incorporating functional data into the neuronavigational datasets. Modern imaging data that might have not been considered for planning purposes in the past, such as fMRI, is gaining significance [BHWB07]. When designing planning tools, these developments have to be accounted for and scalability must be guaranteed. Some of the existing systems limit themselves by design, others are limited by their system architecture with regards to future scalability (e.g., by the reliance on video memory, which is limited in typical consumer hardware).

In this thesis we propose BRAINXPLORE, an interactive data exploration and filtering system to support the neurosurgeon in planning biopsies. BRAINXPLORE offers a set of filtering operations that allow the surgeon to incrementally refine a set of potential trajectories until an optimum has been found.

We solve the issue of scalability by including a **spatial index server** that has previously successfully been used in the research of fly brains [BŠG⁺09]. This component in our work-flow allows us, in contrast to existing solutions, to access a very high number of structures at an arbitrary spatial resolution. Šoltészová [Šol10] has presented the efficiency of the index structure for large distance tables of 2GB.

Besides supporting the surgeon, we also aim at creating insight when it comes to the decision process. To evaluate our system, we had a neurosurgeon use BRAINXPLORE to go through the planning process on a real-world data set of a past biopsy.

The remainder of this thesis is structured as follows: **Chapter 1**, provides information of the methods used in medical image acquisition and data segmentation. An overview over the planning guidelines as laid down by a domain expert at our partner hospital

is given. Finally, we present a real case study to illustrate the state of the art in brain biopsy planning. **Chapter 2** evaluates related work in the fields of biopsy planning and visualization of multivariate data. **Chapter 4** offers an overview of the software and hardware we used for the implementation. **Chapter 5** provides information on the methods we have employed when designing and implementing BRAINXPLORE. Implementation details are outlined in **Chapter 6**. **Chapter 7** evaluates our system based on a real data set. Finally, **Chapter 8** summarizes the contributions of this thesis and proposes future work.

Medical Background

Stereotactic brain biopsies are an efficient and safe [TFM⁺09] method of obtaining tissue samples of brain tumors for further histological evaluation. The planning process builds upon detailed knowledge of the human brain. Because no two brains are identical, medical images have to be obtained for every patient. Orringer et al. [OGJ12] state that "currently, all systems for stereotactic biopsy rely on navigation based on preoperatively obtained images". To build a (semi)automated planning system, structures at risk need to be segmented from the obtained data. Blood vessels are arguably the most critical of these structures. Another commonly considered example is the ventricular system, comprised of the lateral ventricles, third and fourth ventricle.

This Chapter is structured as follows. In Section 2.1, we describe the medical imaging methods that are employed to acquire the data needed to plan a biopsy. Section 2.2 describes the process of segmenting structures at risk from the obtained data. Finally, planning guidelines from related literature are presented in Section 2.3.

2.1 Medical Imaging Methods

Many anatomical structures can be considered in the process of planning a brain biopsy. In order to use data, it first needs to be acquired. The following sections describe the data acquisition methods and explain the relevance of the different data modalities to the planning process.

2.1.1 Computed tomography

As presented in Chapter 3, in brain biopsy planning, the skull is often used to automatically calculate candidate entry points. The best modality for skull image data acquisition is **computed tomography** (CT).

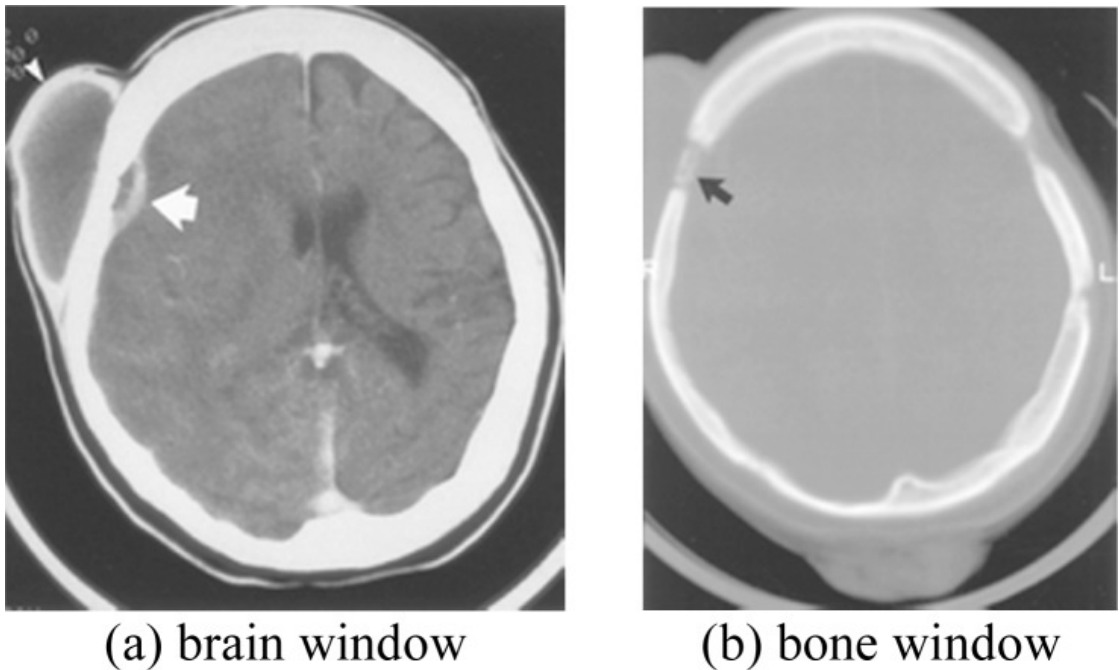


Figure 2.1: CT windowing as illustrated by Xue et al. [XAL⁺12]. (a): Brain window. (b): Bone window.

CT, specifically X-ray CT, is an imaging method that combines X-rays taken from different angles to produce a 3D model of the respective area. Each voxel gets assigned a value corresponding to its radio density μ .

CT offers a standardized unit for describing radio densities, namely, the Hounsfield scale (HU). Air has a value of -1000HU (μ_{air}), water is 0HU (μ_{water}). All other values are scales with regards to the radio density of air and water, as described in Equation 2.1. Depending on the structures of interest, CT data can be windowed in order to optimize contrast. Windowing describes the process of limiting the relevant scale of values to a defined area, i.e., the window. A window is defined by a center and a width. Grey values are then mapped to the values that fall into the interval $[center - \frac{width}{2}; center + \frac{width}{2}]$, with the value $center - \frac{width}{2}$ being mapped to black and $center + \frac{width}{2}$ being mapped to white. Table 2.1 gives examples of commonly used windows. Figure 2.1 illustrates the same dataset, displayed with a bone window and a brain window.

$$HU = 1000 * \frac{\mu - \mu_{water}}{\mu_{water} - \mu_{air}} \quad (2.1)$$

Due to the use of ionizing radiation, CT comes at a risk. Based on their risk estimate from 1991 through 1996, Brenner et al. state that "it has been estimated that about 0.4% of all cancers in the United States may be attributable to the radiation from CT studies" [BH07] and go on to project that "this estimate might now be in the range of

	Center	Width
Lung	-650	1500
Emphysema	-800	800
Soft tissue without contrast agent	40	400
Liver without contrast agent	40	200
Soft tissue with contrast agent	70	400
Liver with contrast agent	60 - 100	300
Neck with contrast agent	50	300
CT Angiography	100 - 200	500
Bones	500	2000
Osteoporosis	300	1000 - 1500

Table 2.1: Commonly used CT windows, adapted from Prokop et al. [PGSP06]

1.5 to 2.0%" based on the current CT use. Hence, eliminating CT scans from the biopsy planning process does not only make data fusion simpler and the whole process cheaper (by eliminating an additional study), but also potentially preserves patient health.

2.1.2 Magnetic Resonance Imaging

In biopsy planning, differentiating between different tissue types, e.g., in the brain, is crucial. For soft tissues, **Magnetic Resonance Imaging** (MRI) is better suited than CT, as it provides much more detail.

MRI is based on the relaxation of hydrogen atoms after excitation by a strong electromagnetic field. Schild [Sch12] describes the process quite simply: "1) the patient is placed in a magnet, 2) radio wave is sent in, 3) the radio wave is turned off, 4) the patient emits a signal, which is received and used for 5) reconstruction of the picture" [Sch12].

Atoms consist of a nucleus and a shell. Inside the nucleus, amongst other things, are protons, i.e., positively charged particles. Protons spin around their axes and thereby create an electrical current. While protons are normally aligned in a random fashion, exposure to a strong magnetic field aligns them either parallel or anti-parallel to the external magnetic field. Slightly more protons are aligned parallel than anti-parallel. Protons precess around the magnetic fields longitudinal axis. The precession frequency is dependent of the magnetic field strength and can be calculated as stated in Equation 2.2. ω_0 specifies the precession frequency, γ is the gyro-magnetic ratio and B_0 indicates the magnetic field strength.

$$\omega_0 = \gamma B_0 \tag{2.2}$$

Parallely and anti-parallelly aligned protons that point in opposite directions cancel each other out when it comes to the net electric current they create. Since there are more parallelly than anti-parallelly aligned protons, however, not all current is

eliminated. When considering the reduced set of protons, i.e., only those with a parallel alignment, elimination occurs in the plane perpendicular to the magnetic field. Let Z be the axis defined by the magnetic field. Then, currents in the XY plane that point to opposite directions cancel each other out. In contrast, all parallelly aligned protons induce magnetic vectors that point in the same direction along the Z axis and hence add up. This magnetization is referred to as longitudinal magnetization. Because the magnetization is longitudinal to the external magnetic field, however, it can not directly be measured.

To resolve this issue, a radio frequency (RF) pulse is applied to the magnetic field. The frequency is based on the Larmor Equation. By matching the radio frequency to the protons precession frequency, resonance is created and energy transfer occurs. For one, this energy transfer flips the alignment of some protons, such that parallelly aligned protons take an anti-parallel alignment and hence, the longitudinal magnetization is reduced. On the other hand, the RF pulse puts the protons in sync, i.e., they all point to the same direction in the XY plane at the same time. The resulting magnetic vector is referred to as transverse magnetization. Due to the precession of the protons, the magnetic vector is constantly moving, thus inducing an electric current that also has the precession frequency. This current is the MR signal. After turning off the RF pulse, the protons whose alignment has been flipped from parallel to anti-parallel take their initial alignment again, thus increasing the longitudinal magnetization. This process is known as longitudinal relaxation or spin-lattice relaxation (as the protons hand over energy to their surroundings, the so-called lattice). The relaxation is described by a time constant T_1 . Conversely, the transverse relaxation (also called spin-spin relaxation), is described by a time constant T_2 . Different tissues have different relaxation times. While water, for example, has a long T_1 and long T_2 , fat has a short T_1 and a shorter T_2 than water [Sch12].

Various sequences can be used to weigh an MRI. The most common are T1-weighted, T2-weighted, T1-weighted with contrast, Diffusion and FLAIR. T1-weighted MRI is particularly useful for assessing the cerebral cortex. T2-weighted MRI helps in detecting edema and inflammation and revealing white matter lesions. Diffusion weighted imaging (DWI) can be used to perform tractography and identify neural tracts.

While the Hounsfield scale used for CT offers a standardized scale, MRI images taken for the same patient on the same scanner at different times may appear different from each other due to a variety of scanner-dependent variations [NUZ00]. This means that absolute intensity values have no fixed meaning and windowing has to be performed for each image. More recently, variants of data standardization have been proposed, such as the work performed by Nyul et al. [NUZ00].

Contrast agents can be used if necessary. Especially in the context of tumor detection, they further enhance the image contrast for more accurate cancer detection and diagnosis [ZL13]. Paramagnetic compounds such as gadolinium based agents are the most commonly used contrast agents in clinical MRI [ZL13]. They mainly reduce the longitudinal (T_1) relaxation property and result in a brighter signal [Sho13, ZL13]. Super-




Figure 2.2: T1-weighted, contrast enhanced MRI. Left Upper: 3D reconstruction. Right Upper: Axial View. Left lower: Sagittal View. Right Lower: Coronal View

paramagnetic nanoparticles (SPMNPs) such as iron oxides, on the other hand, have a strong effect on the transverse (T2) relaxation properties. Figure 2.2 shows an example of a T1-weighted MRI after gadolinium injection in coronal, sagittal, and axial views as well as a 3D reconstruction.

In contrast to CT, MRI has no ionizing radiation and is therefore deemed safer for most applications. On the other hand, acquiring MRI takes a lot longer than CT. MRI is better suited than CT for evaluating soft tissue structures. For evaluating bony structures, CT is more suitable.

2.1.3 Diffusion-weighted Magnetic Resonance Imaging

White matter tracts, such as the **corticospinal tract** (CST) must not be damaged during neuro surgeries, as this could potentially lead to lasting impairments of the patient. One-sided damage to the CST, for example, usually leads to a paresis of the contralateral side. While it is known that most of the CST's neurons originate in the primary motor cortex or the premotor frontal areas, the precise location must be found for each patient individually before planning a biopsy. However, brain tracts are not identifiable by a direct exam in MRI.

Diffusion, in its basic form, is the random movement of molecules. Given a concentration gradient, i.e., a separation between areas of high and low concentrations of molecules, this random movement will result in molecules passing from the area of higher concentration to the area of lower concentration (and vice versa, but probability dictates that the amount of movement from high to low will be higher). This phenomenon is called **flux**

(J) and is defined by Fick's first law as presented in Equation 2.3. The diffusion coefficient D equates the flow vector and the concentration gradient c .


$$J = -D\nabla c \quad (2.3)$$

Diffusion-weighted Magnetic Resonance Imaging (DWI) is an imaging method that builds on MRI and exploits the diffusion patterns of water molecules to reveal microscopic details about tissues. DWI provides an estimate of the **rate of water diffusion** for each voxel. This information can be used to detect early stages of a pathological change.

To acquire a DWI image, first an RF is applied to the subject as with standard MRI, to excite water molecules. Then, a positive, so-called *dephasing* gradient magnetic field is set up. This dephasing gradient will cause electrons in water atoms to spin at different speeds with respect to their spatial location inside the gradient field (they are brought out of phase, hence, dephasing gradient). Once the gradient is turned off (i.e., a homogeneous magnetic field is restored), rotation speed will be the same again for all molecules. Applying a gradient inverse to the initial one will again cause rotation speeds to change with respect to the corresponding atoms location. If the same field strength was applied for the same duration, all molecules should rotate at the same speed and be in phase again. Hence, this negative gradient is called the *refocusing* gradient. In a static system, the refocusing would be perfect. However, if water atoms move due to diffusion, this process will be imperfect, i.e., after the refocusing gradient was applied, not all the molecules will be in phase. This imperfect refocusing will result in a signal loss, which can be used to measure the diffusion constant. However, only the diffusion along the gradient direction can be measured. To obtain Fiber directions, multiple scans are necessary.

Diffusion Tensor Imaging (DTI) expands on DWI by considering not only the rate, but also the **average direction** of diffusion. In areas that are made up of an internal fibrous structure, water will diffuse faster in the direction parallel to the structure [HJM⁺06]. Since a defining characteristic of neuronal tissue is its fibrillar structure, this consideration allows for the identification of brain tracts. As pointed out, this is not possible in T1-weighted or T2-weighted MRI. In contrast to DWI, DTI is usually built on six or more gradient directions, rather than three. Multiple gradient directions (mathematically, at least six, although often more are used) are necessary to calculate a diffusion tensor.

In anisotropic media, the diffusion coefficient D depends on the direction. Therefore, in DTI, it is replaced by the diffusion tensor \mathbf{D} . \mathbf{D} is a 3 x 3 matrix that fully characterizes diffusion in 3D space [HJM⁺06] and can be calculated with the help of the Stejskal-Tanner-Equation presented in Equation 2.4. g denotes the gradient direction. The b -value specifies the diffusion weighting and is proportional to the product of the square of the gradient strength q and the diffusion time interval Δ [HJM⁺06] as illustrated in Equation 2.5. S_0 is the signal intensity without the diffusion weighting, S is the signal



(a) 3D visualization of white matter tracts in MeVisLab. (b) 3D visualization of white matter tracts in MeVisLab.

Figure 2.3: Visualization of DTI data in 2D (a) and 3D (b). XYZ diffusion directions are color coded in RGB. Red fibers indicate that diffusion primarily takes place in the X-direction, green shows diffusion in the Y-direction and blue in the Z-direction. Other colors indicate intermediate diffusion directions.

with the gradient. Since D possesses six degrees of freedom, at least six directional gradients must be applied to fully determine the system.

$$S(g) = S_0 * e^{-bg^T Dg} \quad (2.4)$$

$$b \approx q^2 \Delta \quad (2.5)$$

The diffusion tensor gives, for each voxel, information on the diffusion direction. Usually, directions are RGB encoded, as illustrated in Figure 2.3. In a similar fashion, blood flow across different brain regions can be identified from **functional MRI** (fMRI). This modality can be used to segment functional brain areas.

2.1.4 Positron Emission Tomography

Positron Emission Tomography (PET) and **single photon emission computed tomography** (SPECT) are imaging methods that offer insight into functional metabolic processes. As the name implies, it is based on the observation of positron annihilation events. A radioactive tracer is injected into the bloodstream. Usually, this tracer is glucose-based, as regional glucose uptake corresponds to a tissue's metabolic activity. If a radioactive (i.e., unstable) atom transforms into a stable atom by emitting either a positron or an electron, this process is called beta-decay. Specifically, in the case of a positron and a neutron being emitted, it is called a β^+ -decay. Equation 2.6 [VBTM03] illustrates the process, with p denoting a proton, n being a neutron, e^+ being an electron and finally, ν being a neutrino.

$$p \rightarrow n + e^+ + \nu \quad (2.6)$$

As the tracer undergoes β^+ -decay, it emits a positron which, eventually, interacts with an electron. Since the positron has the opposite charge of the electron, both are annihilated during the interaction. The annihilation results in a pair of annihilation photons, i.e., γ -quanta, moving at an angle of 180 degree to each other. On their way, those quanta interact with their surrounding tissues and are attenuated. Once they leave the body, the photons are measured by means of a scintillator. Based on their location and attenuation, the physician can draw conclusions regarding the spatial distribution of the tracer. PET images are functional rather than topological imaging techniques and hence results have to be correlated to other data, e.g., CT scans. Modern CT devices are capable of acquiring both image modalities in the same session.

In this context, the most important application of PET scans is the examination of tumors with regards to metastases. Since the cost of PET scans is fairly high, no such data has been considered in the case studies presented in this work.

2.2 Brain Anatomy Segmentation

Segmenting the brain at a structural level is not trivial. Figure 2.4 shows intensity histograms for different subcortical brain structures. Based on their intensity values alone, the Amygdala (Am) and the Hippocampus (Hp) are virtually indistinguishable. However, by including spatial information in the segmentation process (i.e., the Amygdala residing to the anterior and superior to the Hippocampus), the inherent ambiguity of class intensity distributions can be overcome. To incorporate this spatial information, brain segmentation is facilitated by the construction of a probabilistic atlas. In such an atlas, somewhat arbitrary raw image coordinates are transformed to a space in which coordinates have an anatomical meaning. Atlases can be based on single specimens, such as the Talairach atlas, or can be constructed from a large population, as presented by Mazziotta et al. [MTE⁺95].

In the context of brain segmentation, different functional levels of increasing granularity can be observed, as illustrated in Figure 2.5. At the topmost level is the hemisphere. The **Hemisphere** is important because crossing the mid-sagittal plane should be avoided in brain biopsies. Beneath the Hemisphere Level lies the **Lobe** Level. Here, the four main Lobes (Frontal, Temporal, Parietal, and Occipital), the Limbic Lobe and a sub-lobar region can be identified. Further down in the hierarchy is the **Gyrus** Level. In the Gyrus Level, anatomical classes (e.g., Caudate, Thalamus, Lentiform, etc.) are defined. In contrast, the **Tissue** Level defines the type of tissue within a voxel (i.e., white matter, grey matter, cerebrospinal fluid). At the lowest level, the **Cell** Level, the brain can be labeled by cell-type in the cortex using Brodmann's scheme [Gar94], presented in Figure 2.6. This is important as certain cortical areas must not be damaged, as explained in Section 4.1.1.




Figure 2.4: Intensity histograms for different brain structures as presented by Fischl et al. [FSB⁺02]. While structures such as the Putamen (Pu) and Thalamus (Th) have significantly different T_1 properties, the Amygdala (Am) and the Hippocampus (Hp) are virtually indistinguishable.

This hierarchical structure narrows down the number of candidate structures for each given atlas position, as the spatial relations between the structures are known. In fact, once the MRI coordinates are transformed into probabilistic atlas coordinates, "the number of classes at a given location is rarely greater than 4, and in fact averages less than 3 within the brain", as stated by Fischl et al. [FSB⁺02]. A probabilistic atlas is a reference system of the human brain that takes the anatomic variability between individual human brains into account.

The Talairach atlas is an example of a non-probabilistic atlas. It was derived from a single sample, namely the brain of a 60 year old French woman. The hemispheres are assumed to be symmetrical. For areas of low intersubject variability and for sites close to the landmarks of the reference system, this type of atlas can be employed successfully [MTE⁺95]. In areas of high intersubject variability and in brain regions known to be highly asymmetric between the two hemispheres, this type of atlas performs progressively less accurate, though.

Conversely, a probabilistic atlas is based on a population of multiple brains. Mazziotta et al. [MTE⁺95] state that "by its very nature, a probabilistic atlas improves in accuracy with time, achieving better statistics as more information is added".

Different packages exist that are capable of performing an **automated parcellation**




Figure 2.5: The five hierarchical levels in the Talairach Atlas as described by Lancaster et al. [LWP⁺00]: Hemisphere Level, Lobe Level, Gyrus Level, Tissue Level, Cell Level

of subcortical structures. Two commonly used packages are **FreeSurfer** [RSRF12] and **FIRST** [PSKJ11]. **FIRST** is a part of **FMRIB's Software Library (FSL)** [SJW⁺04]. Studies [MPX⁺09] have compared the performance of those packages.

Morey et al. [MPX⁺09] used both packages to automatically segment the Amygdala and Hippocampus and compared the respective results to hand tracings performed by a single expert rater with the experience of performing over 500 hippocampal tracings. The authors conclude that FSL and FreeSurfer are not equal when compared to manual tracing and further state that FreeSurfer is generally preferred to FIRST. Conversely, Perlaki et al. [PHN⁺17] find that FIRST is superior to FreeSurfer with regards to putaminal segmentation.

Table 2.3 gives an overview of each packages strengths and weaknesses. FreeSurfer is capable of segmenting more structures than FSL. Table 2.2 gives an overview over each packages set of subcortical labels.

Cortical Parcellation is based on standard atlases, as described above. Both FreeSurfer and FSL offer different atlases. For the cortical parcellation used in the data segmentation process of our case study (see Section 7), the **Destrieux atlas** [DFDH10] was used. It was built on a training set of twelve data sets and segments the cortex into 74 regions of interest per hemisphere.




Figure 2.6: Brodmann's areas. The primary somatosensory cortex, motor cortex and Broca's area, amongst other cortical structures, must not be damaged during the biopsy.

2.3 Planning Guidelines

Quick-Weller et al. [QWKD⁺] state that while stereotactic brain biopsies can be carried out under local or general anesthesia, stress levels appear to be higher in patients undergoing local anesthesia [QWKD⁺].

Certain guidelines have to be adhered to when performing a brain biopsy. The first step in planning a brain biopsy consists of the selection of an appropriate target point. Germano and Kondo [ABB⁺02] recommend choosing target points at the tumors periphery. This is because the center of the lesion may consist of fluid or necrotic tissue. In the latter case, the obtained sample would not be diagnostic. In the former case, drainage of the fluid causes brain shifting, preventing additional biopsies with the same setting.

After a target point was identified, a suitable entry point must be found. Those two points define the biopsy trajectory, which intersects the skull at an entry angle α . Zombori et al. [ZRN⁺14] point out that "from a surgical point of view the entry angle has to be as close to perpendicular as possible, otherwise it is not possible to drill the borehole through the skull".

When deciding for or against a trajectory, caution regarding the vascular system is of the utmost importance. A certain safety margin to nearby vessels needs to be respected, as damaging a brain vessel may lead to irreparable damage or even the death of the patient. Care must also be taken not to damage eloquent areas around the tumor, such

2. MEDICAL BACKGROUND

Structure	FSL / FIRST	FreeSurfer
Left / Right Thalamus	✓	✓
Left / Right Caudate	✓	✓
Left / Right Putamen	✓	✓
Left / Right Pallidum	✓	✓
Left / Right Hippocampus	✓	✓
Left / Right Amygdala	✓	✓
Left / Right Accumbens Area	✓	✓
Brain-Stem /4th Ventricle	✓	✓
Left / Right Thalamus	✗	✓
Left / Right Caudate	✗	✓
Left / Right Putamen	✗	✓
Left / Right Pallidum	✗	✓
Left / Right Cerebellum Cortex	✗	✓
Left / Right Cerebellum White Matter	✗	✓
Left / Right Cerebral Cortex	✗	✓
Left / Right Cerebral White Matter	✗	✓
Left / Right Ventral Diencephalon	✗	✓
Lesion	✗	✓
Third Ventricle	✗	✓
Lateral Ventricle	✗	✓
Inferior Lateral Ventricle	✗	✓

Table 2.2: Subcortical structures segmented by FSL and FreeSurfer, respectively. In the FSL Model, the fourth ventricle is combined with the Brain-Stem, and so the labels for the Brain-Stem also cover the fourth ventricle.

as Broca’s area and motor and sensory cortex [ABB⁺02]. Another safety criteria regards the cortical insertion site. Essert et al. [EHJ10] state that "Vessels are numerous in the brain, and generally located in cortical sulci. Unfortunately they are often invisible when images are acquired without contrast agent or Angiography". Hence, intersecting sulci poses a threat and should be avoided. Trans-gyral trajectories are therefore preferred over trans-sulcal trajectories.

In an interview, our domain expert further pointed out that he discourages puncturing the mid-sagittal plane. Hence, needle entry points should be located on the same hemisphere as the tumor target point. The ventricular system consists of four ventricles: the **lateral ventricles**, the third ventricle and the fourth ventricles. The ventricles are linked to the production of cerebrospinal fluid (CSF). If possible, ventricles should not be intersected, as the puncture of ventricles significantly correlates with postoperative altered mental status and subsequent increased length of hospital stay [GBHM⁺11]. Also, for cosmetics reasons, trajectory entry points above the hairline are preferred in order to avoid visible operation scars.

Attribute	FreeSurfer	FSL / FIRST	Source(s)
Number of segmented tissues	37	16	
Hippocampal Volume Overlap	Higher	Lower	[MPX ⁺ 09]
Hippocampal Volume Difference	Lower	Higher	[MPX ⁺ 09]
Hippocampal Sample Size Estimates	Closer	Further	[MPX ⁺ 09]
Hippocampal Correlation to Hand Tracing	Higher	Lower	[MPX ⁺ 09, SBH ⁺ 16]
Hippocampal 3D Shape Analysis	More Accurate	Less Accurate	[MPX ⁺ 09]
Amygdala Volume Overlap	Equal	Equal	[MPX ⁺ 09]
Amygdala Volume Difference	Higher	Lower	[MPX ⁺ 09]
Amygdala Sample Size Estimates	Equal	Equal	[MPX ⁺ 09]
Amygdala Correlation to Hand Tracing	Higher	Lower	[MPX ⁺ 09, SBH ⁺ 16]
Amygdala 3D Shape Analysis	Less Accurate	More Accurate	[MPX ⁺ 09]

Table 2.3: A comparison of FSL / FIRST and FreeSurfer.

After the biopsy was performed, post-operative CT scans can be taken to exclude the presence of hemorrhages [Hal98] and confirm the location of the biopsy site [TFM⁺09]. Fritsch et al. [FLG⁺98] state that if the post-operative images show no complications, it appears safe to discharge patients the same day or within 24 hours after the biopsy.

Related Work

In the previous chapter, the medical background of brain biopsies was presented and established planning guidelines were explained. The aim of a biopsy planning system is to incorporate medical data and aid the neurosurgeon in finding a suitable needle trajectory. Such a system must provide overview of the medical data in a concise and understandable manner, as well as implement the planning guidelines which are applied in the respective hospital. In this chapter, relevant state-of-the-art solutions in the field of brain biopsy planning on the one hand and medical information visualization on the other hand are summarized.

This Chapter is structured as follows. First, existing solutions for brain biopsy planning are presented in Section 3.1. Second, methods of medical information visualization are shown in Section 3.2.

3.1 Biopsy Planning

In this section, we review the state-of-the-art solutions for brain biopsy planning. In Section 3.1 will give an overview of current brain biopsy planning systems. Section 3.1.1 offers information on the different **work-flows** implemented in the existing systems. **Data acquisition and Segmentation** is outlined in Section 3.1.2. The **Rules and Guidelines** for planning trajectories are presented in Section 3.1.3. **Visualization** modalities for calculated trajectories are shown in Section 3.1.4. Finally, in Section 3.1.5, the existing solutions are summarized and their **limitations**, along with the **improvements** introduced by BRAINXPLORE, are discussed.

3.1.1 Work-Flow

Automatic and semi-automatic trajectory planning is a research topic that occurs mainly in Brain Biopsy (BP) Planning [KKMS11, HMG11, HMP⁺12, GCH⁺14] and DBS plan-




Figure 3.1: Work-flow as presented by Brunenberg et al. [BVVV⁺07]

ning [BLE⁺13, GPP07, EHJ10, ZRN⁺14, ZBM⁺13, BVVV⁺07, BDS⁺13]. It aims at reducing the time surgeons need to spend at planning operations. Hence, a relatively simple work-flow including a high degree of automation is beneficial. On the other hand, as the surgeon is ultimately responsible for the operation’s outcome and hence, his patient’s well-being, certain decisions can not be taken away from him.

Berriault et al. [BDS⁺13] illustrate the fundamental work-flow of automatic trajectory planning in four subsequent steps. The first of these is MRI acquisition or, more generally, data acquisition. In a second step, critical structures are segmented. This can be achieved with or without user interaction. Following this segmentation, target identification makes up the third step. Here, a target for the electrode / needle is specified. Finally, trajectories are automatically calculated and analyzed. From a list of best trajectories, the surgeon can pick one.

Brunenberg et al. [BVVV⁺07] employ a two-pass approach by evaluating MR data first and only integrating CT data in a second step. On the MR side, contrast-enhanced T1-weighted and T2-weighted images are acquired some days before the operation. These images are then processed in three subsequent steps. First, the images are registered and an indirect atlas-based segmentation is performed in order to locate the subthalamic nucleus. Second, an automatic segmentation of critical structures is done. Lastly, the target points and segmentation results are used to calculate possible paths. On the operation day, the integration of CT data follows. First, the CT images are registered to the MR images. Second, the possible paths are visualized in 3D. Figure 3.1 illustrates the work-flow.




Figure 3.2: Segmentation result as presented in [ZBM⁺13]. Gray matter, Hippocampus and Amygdala segmentation are based on T1-weighted MRI before gadolinium injection ((A) and (B)). MRI taken post gadolinium injection is used for vessel segmentation (C). (D) is for evaluation purposes only.

Bock et al. [BLE⁺13] include additional data, namely, the results of microelectrode recording (MER). MER is an intra-operative data acquisition method where electrodes are inserted into the access path. There, the electrodes measure the brain's electrical field and hence, allow for the identification of functional brain regions. In order to make efficient use of this additional data, the authors introduce a recording phase to their work-flow. During the MRI based planning phase, a trajectory is defined by the neurosurgeon. Subsequently, during the recording phase, electrodes are inserted into the access path and advanced towards the target region. Functional brain areas are identified by their MER signal and visualized as described in Section 3.1.4. Finally, stimulation electrodes are placed during the placement phase.

3.1.2 Data Acquisition and Segmentation

Proper segmentation is crucial for safe trajectory planning. Here, the segmentation of blood vessels is of utmost importance, as a puncture can lead to severe trauma or death. Furthermore, any path needs to be chosen in such a manner that no cortex sulci be damaged.

In essence, all considered works share the same concept with regards to segmentation. Differences can be observed in the level of segmentation (i.e., number of different structures that are considered) and the choice of external libraries and tools. The general pipeline, however, can be summarized by the tasks **pre-processing**, **brain-segmentation**, **tumor segmentation**, **vessel segmentation**, **gray matter segmentation**, and **white matter segmentation**.

Pre-processing aims at making data easier to segment in subsequent steps. Brunenberg et al. [BVVV⁺07] use coherence-enhancing diffusion in order to decrease the image noise while preserving edges.

Brain segmentation, also called skull stripping, means the removal of the skull and other non-brain tissues like the eyes from the dataset. Essert et al. [EHJ10] segment the brain from a pre-operative, T1-weighted MRI using BrainVISA [CMG⁺01]

Tumor segmentation is necessary in order to define a target region. This can be done automatically or manually. Herghelegiu et al. [HMG11] segment the tumor in a T1-weighted MRI sequence using the watershed algorithm, which is based on the variation of voxel gradients.

Vessel segmentation is crucial as hitting a vessel is arguably the highest risk one faces when performing a biopsy. Proper segmentation is therefore of the utmost importance. Most authors make use of a vesselness-filter, which enhances voxels within tubular structures. Herghelegiu et al. [HMG11] obtain the blood vessel mask by manual segmentation. Zemann et al. [ZBM⁺13] process a post-gadolinium set of MRI with a vesselness-filter that returns a voxel likelihood of blood vessel presence in the [0;1] range. Figure 3.2(C) shows a corresponding MRI. Brunenberg et al. [BVVV⁺07] also use a vesselness measure, based on the work of Frangi et al. [FNVV98]. Beriault et al. [BDS⁺13] fully automatically extract arteries from time-of-flight (TOF) angiography and veins from susceptibility-weighted imaging (SWI) venography. SWI data is processed with a vesselness-filter. Zombori et al. [ZRN⁺14] extract the vasculature from computed tomography angiography (CTA), 3D Phase Contrast MRI and Time-of-flight (ToF) MR using a custom tool.

Gray Matter Segmentation reveals structures at risk in the gray matter. Those include the ventricles, sulci, gyri and other subcortical structures. Different authors propose different segmentation strategies:

Ventricles must not be punctured during the biopsy. Gemmar et al. [GGF⁺08] preprocess the T1-weighted MRI with an nonlinear anisotropic diffusion (NLAD) filter and then apply a modified 3D region growing algorithm to segment the 3rd ventricle. From the 3rd ventricle's symmetry plane, the mid-sagittal plane (MSP) is found. The anterior (AC) and posterior (PC) commissure of the 3rd ventricle is then identified at the anterior and posterior boundaries of the 3rd ventricle. Brunenberg et al. [BVVV⁺07] segment ventricles with a region-growing method based on the work of Schnack et al. [SPB⁺01] in combination with morphological methods presented by Géraud [Gér98]. Essert et al. [EHJ10] segment ventricles semi-automatically using MITK (Medical Imaging Interaction ToolKit [MNMW09]). Beriault et al. [BDS⁺13] segment the ventricles from the T1-weighted MRI using ANIMAL [CZBE99].

Cortical Structures need to be segmented because trans-gyral trajectories are preferred over trans-sulcal trajectories. This is because small vessels in the sulci are usually not detectable in MRI yet still pose a risk. Brunenberg et al. [BVVV⁺07] use morphological operations (i.e., dilation and erosion) based on the methods of Géraud [Gér98] as well as masking of T2- with results from T1-weighted MR and a connected component analysis in combination with the method of Lohman [Loh98] to segment the sulci and gyri. Essert et al. [EHJ10] automatically segment cortical sulci from the previously extracted brain using an algorithm based on curvature information. Beriault et al. [BDS⁺13] and Zemann et al. [ZBM⁺13] segment the sulci from the T1-weighted MRI using ANIMAL [CZBE99].

Subcortical Structures are at risk in the case of tumors which are situated deeper

in the brain, i.e., further away from the cortex. Beriault et al. [BDS⁺13] segment the Caudate from the T1-weighted MRI using ANIMAL [CZBE99]. Zelmann et al. [ZBM⁺13] use ANIMAL, augmented with a template library and label fusion [CP10] to segment the Amygdala and Hippocampus from T1-weighted MRI. In Figure 3.2 (A) and (B), those structures are highlighted in the respective MRI.

White Matter Segmentation reveals structures at risk in the white matter, primarily white matter tracts (e.g. cortico-spinal tract, optic radiation tract). Zelmann et al. [ZBM⁺13] segment the white matter from the T1-weighted MRI using ANIMAL [CZBE99]. Zombori et al. [ZRN⁺14] point out that they segment white matter tracts from DTI data.

3.1.3 Path Planning

After proper segmentation, the goal of most presented works is **automatic or semi-automatic trajectory planning**. Basically, the planning process is performed either on a mesh/geometry base (e.g., the work performed by Gao et al. [GCH⁺14]) or on an volume base (e.g., the work performed by Shamir et al. [STD⁺10]). Regardless of the chosen approach, the general steps can be summarized as follows.

The first step is the **selection of a target point or area**. This choice is usually based on a prior segmentation, e.g. with a watershed algorithm or manual intervention by the surgeon. After a target point or area was chosen, **potential entry points** must be selected. This can be done either fully automatically or based on a user-specified entry region. Usually, a skull mesh is used for sample point generation. Based on the target point or area and an entry point or area, a list of **potential trajectories** is generated. The list of trajectory candidates is processed and each candidate is assessed for quality based on certain criteria. While different authors propose different evaluation criteria, the consensus is that the minimum distance between the trajectory and the closest vessel is of utmost importance. This **trajectory evaluation** yields the set of results. The trajectories that score best in the assessment are presented to the user in a subsequent step.

Brunenberg et al. [BVVV⁺07], automatically generate a set of safe entry points. These points are chosen such that they begin in front of the motor cortex but behind the hairline. From each of these points, a straight path towards the target is calculated based on the segmentation results. For each such path, two cost functions are stored. Those costs are based on the **minimum distance to a vessel** (function one) and the **minimum distance to a ventricle** (function two). For faster search in the result set, safe paths (i.e., paths that do not cross a safe margin of 5 mm around vessels and/or ventricles) are stored in bins representing the respective distances. This way, good paths (i.e., such paths that exhibit higher distances than defined thresholds) can be proposed without further examination.

Gao et al. [GCH⁺14] present a prototype system for craniofacial surgery. and present experimental results for tumor surgery. In a first step, the authors form a spherical target

area around the target point. Trajectories are evaluated based on visibility attributes of this target area. Five measures are used for quantification of a path's quality: **The Vital Distance (VD)** essentially encompasses the same information as considered by Brunenberg et al. [BVVV⁺07]. **The Insertion Depth** gives a measure of the paths length, with shorter paths being superior to longer ones. **The Visible Size of the Target Area (VSTA)** from the entry point yields a safety margin. The target area is defined as a sphere centered around the target point. Higher visibility means less occlusion and, as a consequence, higher distances to vital structures. Hence, paths with a larger visible size of the target area are preferable. The authors claim that if the VSTA is above a specified threshold, the safety of the path can be guaranteed. **The Absolute Visible Size (AVS)** accounts for the fact that tumors and similar structures are not necessarily spherical in their appearance. Therefore, the AVS is defined as the visibility of the segmented target structure from the starting point, when considering the occlusions through surrounding tissues. **The Relative Visible Rate (RVR)** of the target is the AVS divided by the **whole area size (WAS)** of the target. The WAS is the size of the tumor from the start point of the path when no surrounding tissues are considered. Like the AVS, the RVT gives a measure of occlusion and thus, risk of crossing a vital structure.

The scene is bounded by a sphere, which is sampled with approximately 18k sample points. Using each sample point as a camera position with the focal point set at the target position, the scene is rendered twice, once with surrounding tissues and once without them. From these renderings, the AVS and RVR can be acquired. These two measures, together with the visible size of the target area, can be used to filter and dispose of unsuitable paths before calculating the other markers for the remaining paths. Since the number of remaining paths can still be too large to allow for a feasible selection, all paths are clustered using a mean shift algorithm as proposed by Cheng [Che95]. Clustered paths are further reduced to the safest one. Implicitly, a safety margin is given by the position inside the cluster, as medial paths naturally offer a bigger margin than lateral ones. Additionally, as outlined before, a shorter path is preferable over a longer one. Figure 3.3 illustrates the concept of the clustering approach.

Essert et al. [EHJ10] introduce additional rules for their DBS planning system. Here, a set of eight rules has been developed in cooperation with participating neurosurgeons. The first four of these rules are Boolean (i.e., they must be satisfied for a path to be acceptable). First, the **target point** needs to be placed at the target. Second, the **entry point** must be on the scalp. Third, **path length** is restricted to a maximum of 90 mm. This eliminates insertion points on the other side of the cranium. Finally, **structures at risk**, i.e., vessels and ventricles, are avoided.

The remaining four rules are preferences (i.e., they need to be optimized based on a set of weights). The first rule concerns the **minimization of the path length**, as longer paths are inherently more damaging and dangerous than shorter ones. This is a low-weight rule and can be omitted. The second rule states that **distances to structures at risk** should be maximized. Again, this is a preference rule, as the avoidance of risky structures

already yields a safety margin. The third rule defines a relationship between the target shape and the **orientation of the electrode**. The trajectory axis should be oriented as close as possible with target main axis. Finally, according to the fourth rule, the **tip of the electrode** should be placed as close as possible to the center of the target. This preference rule is more important for DBS than for biopsy planning, as the electrode remains in the brain with DBS.

While the first four of these rules are Boolean (i.e., they must be satisfied for a path to be acceptable), the remaining four are preferences (i.e., they need to be optimized based on a set of weights). Rules one and two, regarding the position of insertion points and target points, are provided with the input data rather than being determined, hence reducing the model's complexity.

The insertion mesh is evaluated on a triangle-by-triangle basis, with triangles that do not satisfy the hard constraints being eliminated from the solution set. Triangles which only partially satisfy the constraints are further subdivided into four smaller triangles which are then re-tested. From the remaining set of insertion points, some homogeneously spaced points are chosen and evaluated with regards to satisfying of the soft constraints. Around the best of these points, a connected components analysis is performed and Nelder-Mead optimization is done from the best candidate. This way, local minima are avoided. Evaluated triangles are rendered with a transfer function that color codes the aggregative cost function for a trajectory. Good entry regions are presented in green, bad entry regions are presented in red, and a range of progressive intermediate colors is used for the other areas. Figure 3.4 illustrates the color map, applied to the insertion mesh.

Herghelegiu et al. [HMG11] follow a slightly different approach. First, they segment the tumor using a watershed algorithm. Furthermore, they manually segment a blood-vessel mask. Then, they apply a distance transformation (DT) to the vessel-segmentation results. This distance transform yields color-coded information about the proximity of vessels for every pixel in a slice. For every **entry point EP** and every tumor point **TP**, a simple check is performed to determine whether the tumor point is reachable, i.e., whether it lies on the tumor's surface and allows for sufficiently deep entry to actually take a sample. Each point satisfying these requirements is added to a list of reachable points. Around each of those reachable tumor points, a disk **D** is created in the **needle pathway's (NP's)** normal plane as a proxy object for cost calculation. For each voxel **V** on the slice, a direction between the **entry point EP** and V is derived. If a point TP belonging to the tumor is found along that direction, the corresponding NP is stored with its cost being set to the nearest distance to a vessel. This distance can be acquired from the previously calculated DT.

Zombori et al. [ZRN⁺14] point out that for reasons of robustness, the entry angle of the trajectory should be as close to 90 degrees as possible. In their MITK based EpiNav system, a system for DBS planning, the authors consider much more information than all the above presented implementations: **White matter tracts** derived from DTI data, **lesions and eloquent cortex** derived from fMRI, **areas of ictal hyperperfusion** derived from SPECT images, **areas of hypometabolism** derived from PET image, ictal

or interictal EEG/MEG sources, **blood vessels** derived from CTA, 3D Phase Contrast MR imaging and in some cases ToF MR and, the **skull surface** derived from CT, CTA or pseudo CT synthesized from an MR scan.

Entry point candidates are automatically created from the vertices of the skull mesh. In a first step, distance and entry angle for each trajectory are checked against user-defined thresholds. Trajectories that do not match the criteria are immediately discarded. This shrinks the search space in a computationally inexpensive way. The innovation that is presented in EpiNav is the inclusion of a full distance profile along the trajectory, presenting distances to critical structures for 256 sample points per trajectory. The trajectory that exhibits the lowest overall risk is chosen and presented to the surgeon.

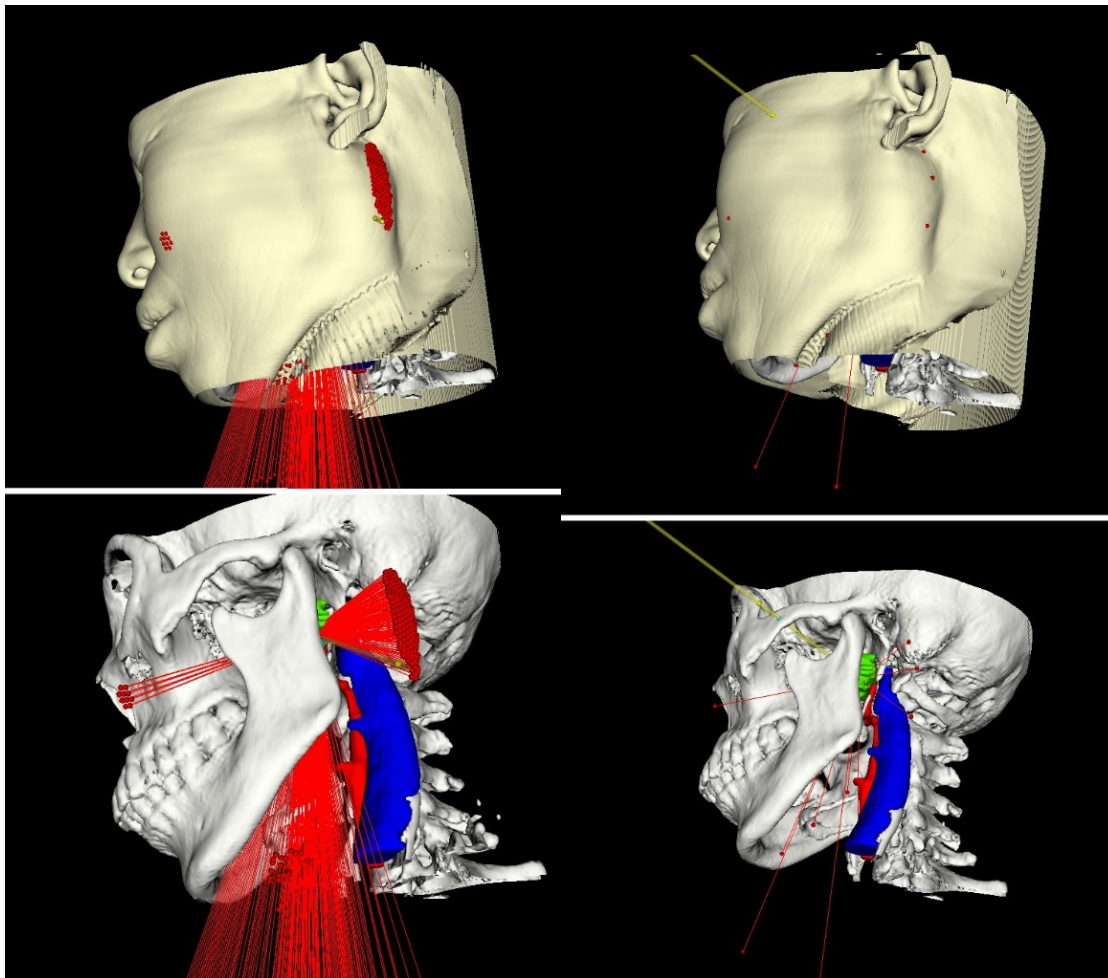
Zelmann et al. [ZBM⁺13] introduce novel scoring factors specific to deep brain stimulation. These factors revolve around the fact that deep electrodes need to cover a large volume for recording purposes. For the purposes of biopsy planning, those factors are not relevant. What is interesting, however, is that the authors chose to pick one thousand, pseudo-randomly chosen (i.e., they are uniformly spaced) trajectories for evaluation. This raises the question whether or not good paths get discarded by chance. Still, validation shows that results are significantly further away from blood vessels ($p < 0.01$) than the manually chosen paths.

3.1.4 Visualization

Although there are obvious differences in the visualization pipelines implemented by the various authors, some shared principles exist. For example, 3D medical images are traditionally displayed in an **MPR view**, i.e., a set of orthogonal cuts through the volume. By selecting a voxel in one of the slices, the linked views displaying the other planes are automatically updated and the planes containing the voxel are presented. More context is offered in **3D views**. Hence, these views are most practical for creating a global understanding of the trajectory. A different use of the MPR view is the **Needle Path Slice View**. To keep the decision transparent, information on the trajectory can be encoded by choosing planes that are orthogonal to the needle path. This way, the volume can be presented from the needle's perspective. This is also called a **Probe's Eye View**. Finally, any number of charts, plots and other well-known **InfoVis** techniques can be employed to communicate certain qualitative or quantitative criteria of a trajectory explicitly. These **Visual Analytics Views** aim at creating insight into data that might not be revealed easily in more traditional views.

Essert et al. [EHJ10] present to the user a global color map which encodes the aggregate cost function for each path. Green pixels indicate the best insertion points whereas red pixels represent the worst points. Since the individual rules are evaluated in distinct color maps, the weighting parameters can be changed interactively, as only the cost aggregation needs to be recalculated.

Gao et al. [GCH⁺14] provide the user with a 3D view illustrating vessels, tumors, and the suggested path. Furthermore, a slice view is given where the path can easily be altered.



(a) Trajectory space before clustering

(b) Trajectory space after clustering

Figure 3.3: Trajectory clustering as presented in the work performed by Gao et al. [GCH⁺14]. Trajectories at the center of a cluster of valid trajectories are considered inherently safer than those at the cluster’s boundary.

For further visualization of the path’s quality, scatterplots and Parallel Coordinates can be displayed.

Likewise, Herghelegiu et al. [HMP⁺12] provide a 3D view along with an augmented 2D slice view (orthogonal to the needle direction), which is special in that at every slice, a line connecting the needle position to the nearest vessel is displayed. Furthermore, a needle pathway distance graph, presenting the nearest distance to a vessel along the trajectory, is given. Figure 3.5 shows a screenshot of the user interface.

In the context of DBS, Bock et al. [BLE⁺13] propose a visualization modality that supports the neurosurgeon by presenting the intersected brain regions as a series of beads,




Figure 3.4: Aggregative cost function evaluation used for path planning as presented by Essert et al. [EHJ10]. Green areas represent good entry points, red areas are considered bad entry points.

as illustrated in Figure 3.6. MER electrodes are inserted along the access path and record the brain's electrical field. From that signal, the functional brain region at the electrodes depth can be identified. For each classified brain region, a bead is attached to the string of beads. The string of beads is rendered behind the electrodes in the application's contextual view. Regions that are outside of the target area result in beads that are rendered in different shades of red, while regions in the target area yield green beads. Finally, crossing uncategorized areas results in a black bead in order to maintain the spatial relationship of a bead string. Although multiple electrodes are employed, only one bead is added for all of them. This is a valid simplification because, as the authors point out, "the different functional regions of the brain are oriented such that either all electrodes detect the same signal (either type of tissue or undefined) or a subset of electrodes detects a regional signal and the others detect an undefined signal".

3.1.5 Limitations

To the best of our knowledge, all existing solutions take into account only a relatively small number of structures. The most basic solutions [HMG11, HMP⁺12] consider only the distance to vessels along the needle trajectory. More sophisticated solutions [BVVV⁺07, EHJ10] consider the distance to vessels, ventricles and cortical sulci. Additionally, Beriault et al. [BDS⁺13] include the caudate. Zemann et al. [ZBM⁺13] consider vessels, sulcus patterns, Amygdala, Hippocampus, as well as white and grey matter in the temporal lobe. Zombori et al. [ZRN⁺14] present the most sophisticated approach, considering vessels, white matter tracts, lesions, eloquent areas, as well as PET and **Single Positron Emission Tomography** (SPECT) data . Since their implementation runs on the GPU, though, GPU memory limitations might be a potential bottleneck. Converting MRI




Figure 3.5: Visualization of the Biopsy Planner presented by Herghelegiu et al. [HMP⁺12]. Spatial information on the needle trajectory can be derived from the 3D and 2D views. Furthermore, the augmented slice view highlights vessels in the needle trajectories proximity. The needle pathway distance graph plots the needle’s distance to the nearest vessel along the trajectory. Finally, the entry points stability map presents the safety margin for each potential entry point inside an ellipsoidal region of interest defined by the user.

data, whose native format is voxel-based, to geometry meshes yields an additional loss of precision. If this loss of precision is to be avoided, a volume-based approach must be taken. In this case, video memory may be too small for a sufficiently high number of structures at risk. A typical MRI resolution is 256^3 . When storing each segmented structure in a 3D texture for further use on the GPU, each texture has a size of 16 MB. At this time, typical consumer hardware offers 2GB of dedicated video memory. This means that 125 distinct structures can be uploaded into the video memory. As brain atlases become more sophisticated, the number of structures at risk to be considered in a biopsy is likely to increase greatly.

We overcome this limitation and allow for an unlimited number of structures at arbitrary resolution by integrating a spatial index server [BŠG⁺09] into BRAINXPLORE that can be queried for distance information in a very fast and efficient manner. The decision to resort to this spatial index instead of performing GPU-based distance computation was made considering that this approach is not limited by GPU memory, but rather offers a convenient way of accessing out-of-core data.

Linked to the limited number of structures is the rigid definition of rules. Systems, which are restricted to a defined number of structures at risk do not necessarily support custom operation guidelines that are unique to a particular hospital. For example, our domain expert noted that in our partner hospital, as a rule, the mid-sagittal plane is not crossed

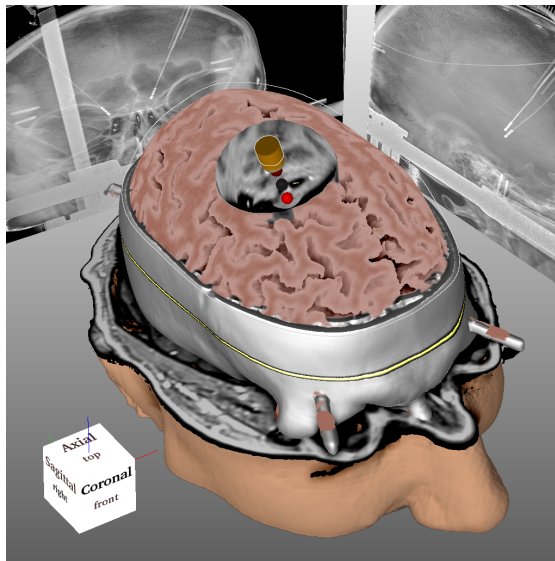


Figure 3.6: Contextual view during the recording phase as proposed by Bock et al. [BLE⁺13] The beads behind the electrode indicate intersected functional brain regions, identified by their MER signal. Red beads are outside of the target area, green beads indicate a region that is within the target area. Black beads represent areas that could not be classified.

during biopsies. To the best of our knowledge, no system in the established body of work allows for such constraints. We offer the neurosurgeon a set of geometric rules (e.g., needle insertion angle at the skull, needle insertion angle at the tumor, trajectory length) as presented in other works [GCH⁺14, EHJ10]. Additionally, we introduce a flexible spatial rule system that can consider any given structure (e.g., the afore-mentioned mid-sagittal plane) and discard trajectories based on their spatial relationship to that structure (e.g., maximal overlap, minimal distance).

3.2 Medical Information Visualization

Information is processed or interpreted data that has a meaning for those who know how to read it. Presenting this information in a meaningful way is the aim of **Information Visualization** (InfoVis). Card et al. define InfoVis as "the use of computer-supported, interactive, visual representation of abstract non-physically based data to amplify cognition" [CMS99]. In the following sections, we review the state of the art of typical InfoVis methods and evaluate their relevance to the interaction techniques we implemented.

In Section 3.2.1, the use of Parallel Coordinates in existing brain biopsy planning systems is presented. Information on different implementations of sparklines and line graphs is given in Section 3.2.2. Finally, in Section 3.2.3, methods for visualizing multivariate data are investigated.

3.2.1 Parallel Coordinates

Parallel Coordinates [Ins85] are a well established InfoVis technique used for displaying multivariate data. Data sets are displayed as lines passing through a set of vertical, parallel axes, where each axis represents a single data feature. By arranging the axes, correlations between data features can be identified.

In the context of brain biopsy planning, Parallel Coordinates have been employed by Gao et al. [GCH⁺14]. The risk signature for each trajectory is represented in the Parallel Coordinate View and updated interactively when the path is adjusted.

Martin and Ward [MW95] evaluated the usefulness of brushing in Parallel Coordinates and found it an "indispensable tool for visualizing and examining" multivariate data. To the best of our knowledge, no existing work in the context of brain biopsy planning implements trajectory picking based on a Parallel Coordinate View.

3.2.2 Line Graphs and Spark Lines

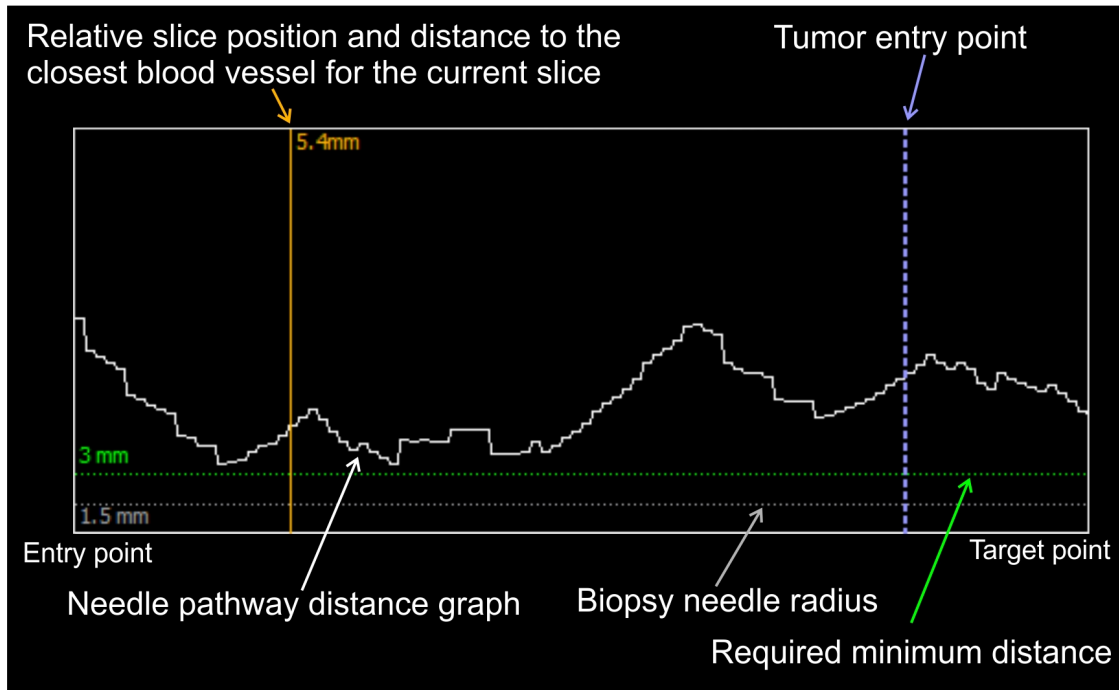
A line graph is a type of chart which, similar to a scatter plot, displays data based on point, so-called markers, on a 2-dimensional grid. Markers are connected by straight line segments. Those graphs offer insight into data trends and help the viewer to understand the relationship between the considered variables. A **spark line** is a special type of line graph that indicates the change over time in some variable. Spark lines are commonly used to illustrate the value of shares in the stock market. In the context of brain biopsy planning, spark lines have been employed to qualitatively visualize the distance between the trajectory and structures at risk.

Herghelegiu et al. [HMP⁺12] present a special form of line graph. A **needle pathway distance graph** shows the distance to the closest blood vessel for every point of the needle pathway. Zombori et al. [ZRN⁺14] expand on this concept by introducing filled areas and annotations indicating the minimum distance to structures, risk thresholds, safety margins and the resulting risk at each sample position. Figure 3.7 illustrates both views.


3.2.3 Methods for Visualizing Multivariate Data

In the context of biopsy planning, the quality of a candidate trajectory depends on a multitude of factors. In some views (e.g., summary and filtering views), multivariate data must be displayed in such a manner so as to instantly reveal significant features to the user to drive the selection process. At the same time, dimensionality must be reduced to such a degree that visual cluttering is avoided.

Ropinski et al. [ROP11] discuss **glyph-based** visualization techniques for spatial multivariate medical data. The authors differentiate between a **pre-attentive** and an **attentive** phase of stimulus processing. During the pre-attentive phase, impulses are perceived in parallel, while during the attentive phase, they are perceived sequentially. The authors identify the basic glyph shape, color, transparency, and texture as well as



(a) Needle pathway distance graph as implemented by Herghelegiu et al. [HMP⁺12]. Needle radius and safety margin are represented with grey and green dotted lines.



(b) Needle pathway distance graph as implemented by Zombori et al. [ZRN⁺14]. The grey area represents the minimal distance to a critical structure, red and green lines indicate the lower and upper safety threshold, respectively. Dark blue areas represent the risk at a given sample position.

Figure 3.7: Distance Graphs used in brain biopsy planning

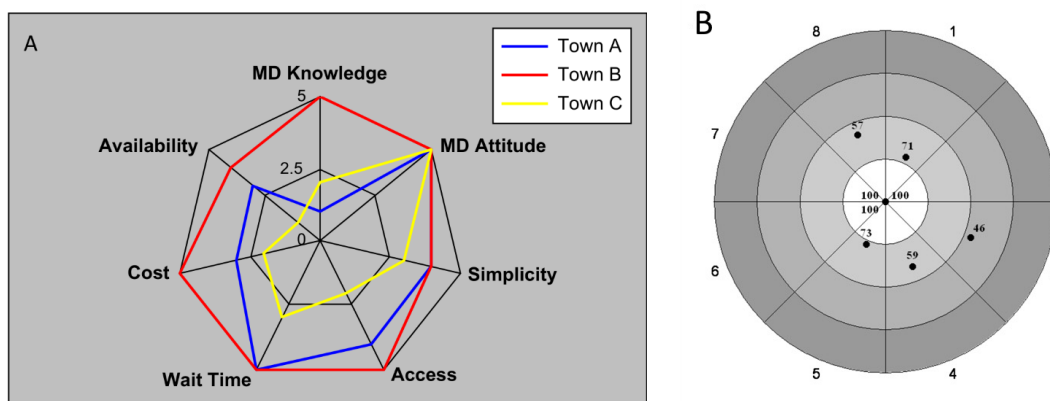


Figure 3.8: (A) Radar chart depicting patient satisfaction across three different towns, as presented by Saary [Saa08]. (B) Target chart depicting the quality of clinical outcomes for a healthcare provider across eight different categories as presented by Stafoggia et al. [SLF⁺11]. The chart immediately conveys an impression of good global performance.

glyph placement as pre-attentive features. Composite glyph shape (e.g., added directional vectors), glyph legend, and interaction details are identified as attentive features.

Stafoggia et al. [SLF⁺11] describe **radar charts**, **target charts** and **spie charts** as means of presenting multivariate data in the context of health care outcomes in an easy-to-convey manner. Radar charts are depictions that map data values to the radii of polygonal coordinate systems. In those coordinate systems, axes (spokes) originate from a shared center (node) and are equidistantly distributed in a circular fashion. Connecting the end points of all axes yields a polygon. On each axis, the corresponding data value is applied and the resulting vertices are connected. The shape of the created polygon offers the viewer a qualitative sense of the data-set's attributes. Saary [Saa08] presents the use of radar charts as a means of comparing multivariate data across different populations. Figure 3.8 (A) illustrates a radar chart that represents the results of an interview where patients were asked to rate their interaction with a hospital according to different criteria. In this case, it is immediately apparent that patients in town B are more happy with their hospital in every regard than patients in the other towns.

Target charts present the user with a pie chart where each circular sector corresponds to one data dimension. Individual data points are mapped to their respective sector, with higher performing data points being drawn closer to and worse performing data points further away from the target center. Figure 3.8 (B) presents a target chart where the quality of clinical outcomes for a single healthcare provider across eight different categories is illustrated. Dots closer to the target center indicate higher scores in the respective categories. Finally, spie charts are superimposed bar charts where the radius of each slice represents the data value and the slice angle represents the feature's relevance.

The authors point out that while the radar chart is easy to implement, "the interpretation

of the results strongly depends on the ordering of the indicators being displayed". Also, missing information introduces ambiguity (i.e., it is impossible to distinguish between a missing data feature and one whose value is zero). They continue to state that the target chart does not perform well in situations where many data dimensions are displayed. Finally, they conclude that spie charts overcome the limitations of both radar charts and target charts and allow for the flexible definition of the different weights of displayed data dimensions. Finally, the authors conclude that out of the three presented techniques, "the spie chart is the best alternative for graphically presenting clinical outcome indicators for comparative evaluations among health care providers or populations".

Chernoff [Che73] proposes **Chernoff Faces**, another form of iconic representation of multivariate data. In this visualization, data features are mapped to 18 distinct facial features on cartoon faces. The author explains that this approach is based on a reverse image recognition algorithm - instead of discriminating between human faces by reducing them to numbers, numbers are discriminated by rendering faces and leaving interpretation to the expert user. Morris et al. [MER00] conclude that Chernoff facial feature perception is not pre-attentive and that hence "Chernoff faces may not have significant advantage over other multivariate iconic visualization techniques".

Visual Computing Work-flow and System Design

After briefly discussing the state-of-the art in brain biopsy planning, in this Chapter we present our prototype for a novel brain biopsy planning framework. It allows the definition of custom planning rules and interactive exploration of the resulting trajectories. Here we explain how the requirements were defined and how our system overcomes the challenges presented by existing solutions.

In Section 4.1, we describe how we conducted the requirement analysis and the results it produced. Section 4.2 offers a global overview of our system and the participating modules. Finally, in Section 4.3, the work-flow from the user's perspective is presented.

4.1 Requirement Analysis

To evaluate the requirements of our new system, we have conducted interviews with the domain expert and asked about the biopsy planning work-flow as performed in the partner hospital. Before deriving requirements, we will discuss the interview in this Section.

Currently, our domain expert performs the biopsy planning with the help of the MedTronic (<http://www.medtronic.com>) StealthStation Surgical Navigation System. In an **orthographic coordinate system**, MRI data is displayed in a sagittal, cranial and coronal view. Also, a **supporting 3D view** is presented to the surgeon. Various data (i.e., CT, MRA, etc) can be used as overlays to highlight structures such as vessels.

To plan the biopsy, as a first step, a **target point** or **area** on the tumor is defined. This target point should be located inside living tissue, as necrotic tissue is not well suited for further histological examination. To define an **entry point**, the surgeon first defines the




Figure 4.1: Cortical regions of interest, adapted from Desikan et al. [DSF+06]. Pial (left) and inflated (right) surfaces of the right hemisphere. Top: Lateral representation of right hemisphere. Bottom: Medial representation of right hemisphere.

suitable quadrant. The mid-sagittal plane should not be crossed. Then, certain areas on the cortical surface are excluded from the set of possible entry points.

The cerebral cortex is divided into the frontal, temporal, parietal and occipital lobe. Each lobe is further subdivided into functional areas. The frontal lobe, for example, is divided into the gyrus frontalis superior (F1), gyrus frontalis medius (F2), and gyrus frontalis inferior (F3). The same partition can be applied to the temporal lobe. It consists of the gyrus temporalis superior (T1), gyrus temporalis medius (T2), and gyrus temporalis inferior (T3). In the parietal lobe, P1 denotes the lobulus parietalis superior, while P2 and P3 together form the lobulus parietalis inferior. Figure 4.1 illustrates the lobes and the pre- and postcentral gyrus, as described in the Desikan [DSF+06] cortical atlas.

Each area has distinct properties that can potentially make it a no-go area. For instance, the area between the pre- and postcentral sulcus (i.e., the central gyrus) contains the motor cortex as well as the somatosensory cortex, as illustrated in Figure 2.6. These areas are highly critical and must not be punctured.

For example, the gyrus frontalis inferior (F3) is located on the left hemisphere. The F3




Figure 4.2: Needle trajectory in the context of subcortical critical structures. Displayed in red is the vascular system. The right and left ventricle are presented in blue and yellow respectively. The tumor is shown in orange. The pyramidal tract is displayed in green.

consists of three parts, the pars orbitalis, pars opercularis, and pars triangularis. The latter two form the so-called Broca's area. Broca's area is linked to motor language and damage to this area can impair speech. The gyrus temporalis superior (T1) contains the primary auditory cortex. Together with the gyrus supramarginalis and gyrus angularis on the left hemisphere, it is linked to sensor language. Damage to these regions might impair the comprehension of language. If possible, the distance between the insertion point on the skull and the insertion point at the cortex should be minimized. Finally, for cosmetic reasons, entry points should be located behind the hairline.

If an entry point has been found that satisfies all the above mentioned criteria, the trajectory is evaluated. Vessels, ventricles and the pyramidal tract, amongst other structures, must not be crossed. Figure 4.2 illustrates a trajectory (blue) in the context of the vascular system, the ventricles and the pyramidal tract.

If a trajectory satisfies all outlined constraints, it is accepted. Otherwise, either the entry or target point are varied and the resulting trajectory is re-evaluated. This process is iteratively repeated until a suitable trajectory has been found. The process is demonstrated in a real-life scenario in Section 4.1.1.

4.1.1 Case Study

To better understand the planning process outlined in Section 4.1, we have carried out a case study on real data from a previous biopsy. Figure 4.3 illustrates the decision finding process. Vessels are highlighted in a light red. The planning process for this case was challenging because of the lesion's proximity to a vessel. The lesion can easily be identified in the MRI as an approximately spherical area of high intensities at the borders and low intensities at the center. The low intensities indicate necrotic tissues which are not suitable for histological analysis. Hence, as a first step in the planning

process, a target point near the lesion's caudal border is defined where the tissue is alive and metabolically active.

In a second step, an entry point is chosen, in accordance with the criteria defined in Section 2.3. Because the target point is located in the left hemisphere and the mid-sagittal plane must not be crossed, the entry point must be on the left side as well. Since the target is also located anterior to the mid-coronal plane, the parietal and occipital lobes are disqualified as entry regions. This leaves the frontal and temporal lobe as candidates. F3 and T1 are excluded due to their anatomical relevance, as explained in Section 2.3. A suitable trajectory was found iteratively, as illustrated in Figures 4.3 and 4.4.

4.1.2 Core Requirements

Based on the state of the art case study presented in Section 4.1.1, we have derived the following **core requirements** for our system:

- Automatic entry point generation
- Multiple target point optimization
- Avoidance of structures at risk

Given the relevant hemisphere, a set of potential entry points must be calculated automatically. These entry points must be on the skull surface and must not be located on areas where trajectories are not feasible such as the eyes or ears. Trajectories must be evaluated for a set of target points rather than a single target. This, in conjunction with the automatic entry point generation, allows the system to evaluate a multitude of candidate trajectories. As presented in the case study, currently the surgeon manually evaluates a relatively small set of trajectories, piece by piece, until a suitable candidate is found. Delegating this task to the system is the main driver for time savings. To ensure the patient's safety, structures at risk such as vessels, ventricles and sulci must not be damaged by any resulting trajectory.

While this list satisfies all requirements that arise from the case study, we have decided to expand it and make some fundamental improvements. First, rather than dictating the selection criteria for valid trajectories, we want our system to allow for a **Flexible Filter Rule Definition**. The system must allow the neurosurgeon to flexibly define rules for trajectory evaluation. Trajectories that violate those rules must not be considered for further evaluation. Furthermore, discriminates valid trajectories with regards to their quality. It also offers the neurosurgeon a way of defining the rules for this **Trajectory Quality Evaluation**.

From an interaction design point of view, the system should allow for **interactive filtering** of trajectories. Initially generated trajectories can be filtered based on attributes such as length, structure to vital distances, etc. Filtering is interactive and provides immediate




Figure 4.3: Red lines indicate rejected trajectories, the green line shows the accepted trajectory. a) The tumor can clearly be identified in the MRI by an elliptic region of higher intensity. The green sphere indicates the chosen target point. The red line shows the first selected trajectory. This path violates the safety distance to the pyramidal tract, visible in the axial view as white area and is therefore invalid. b) The trajectory intersects a ventricle, as seen in the coronal view and is therefore invalid.




Figure 4.4: Red lines indicate rejected trajectories, the green line shows the accepted trajectory. a) This trajectory intersects a sulcus, as seen in the coronal view and is therefore invalid. b) This is the trajectory that was finally accepted.

feedback. Since the neurosurgeon is the authority in his field and ultimately has the responsibility for the biopsies outcome, the system should offer the possibility of manually refining existing trajectories and evaluating the result. The **interactive refinement and addition** of trajectories must be accomplished at interactive framerates.

Because our system is designed to be a proof-of-concept prototype, we have decided to implement different visualization techniques and evaluate them in the context of biopsy planning.

4.1.3 Interaction Requirements

To define requirements for the **user interface** of BRAINXPLORE, we conducted multiple design sessions with a domain expert. We discussed the challenges in his work-flow and utilized tools. Based on our consultation with the expert and previous work, we implemented the following design guidelines for developing the interface:

Workflow-based user interaction: The work-flow, from the user’s perspective, is described in Section 4.3. The user interface should be arranged so that it supports the work-flow rather than hinders it.

Interactivity: Data segmentation is non-trivial and error-prone. Hence, the final decision should always belong to the surgeon. The system should allow to change both endpoints of a suggested trajectory and provide interactive feedback. Also, rule refinement should provide immediate visual feedback.

Providing diverse data perspectives and views: 3D views offer good spatial context but also tend to exhibit a high degree of occlusion. 2D slice views, on the other hand, do not offer the same spatial context but are very good at conveying information on a limited part of the data. In the context of brain biopsy planning, for example, when tracing the needle trajectory, a slice orthogonal to the trajectory can be extracted from the MRI dataset at every point on the trajectory. This slice then offers a direct view of the needle’s surroundings, i.e., a probe’s eye view. Info-Vis views such as Parallel Coordinates offer information that goes beyond spatial information, e.g., data distribution and correlation in multidimensional datasets. For example, histograms can offer an estimate of how a rule refinement may affect the number of available trajectories. Parallel Coordinates offer a convenient picking method that allows the user to select trajectories based on multiple features in one view.

4.2 System Overview

BRAINXPLORE consists of four major components: a data preprocessing module, a trajectory preprocessing module, the index server, and an interactive visualization module (i.e., the user interface). The interactive visualization module is implemented in the rapid prototyping software MeVisLab [Sol17]. Figure 4.5 illustrates our system




Figure 4.5: Overview of BRAINXPLORE system architecture. Data acquisition and segmentation are not part of BRAINXPLORE and can be accomplished using various approaches. The data preprocessing module translates data into BRAINXPLORE’s internal format. The index server stores distance information for all trajectories. The trajectory preprocessing module samples the volume and generates initial trajectories. All user interaction takes place in the user interface module, where rules are defined and refined until a suitable trajectory has been found.

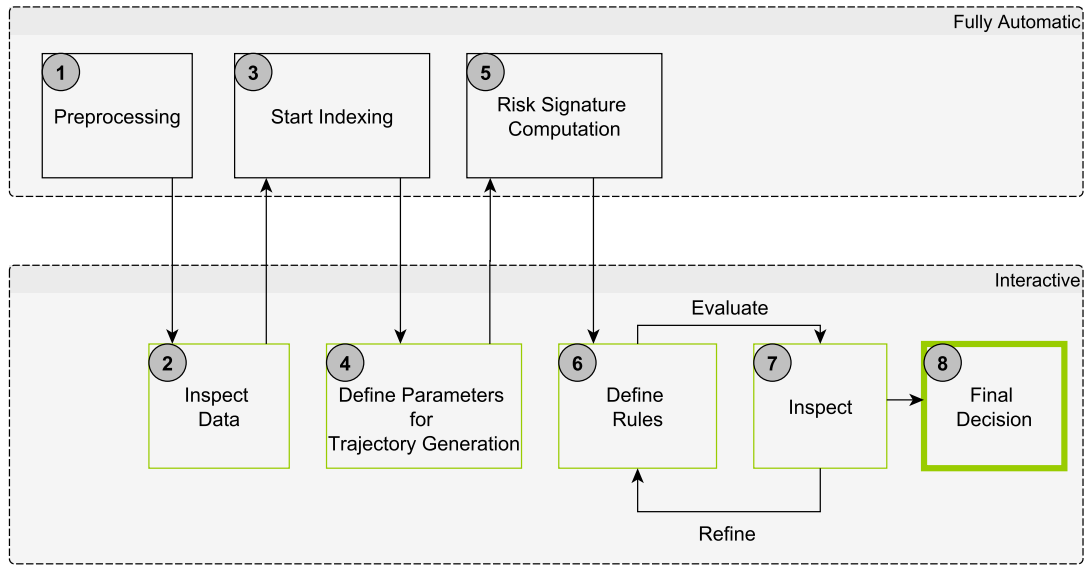


Figure 4.6: User work-flow. Steps that require user interaction are highlighted in green. Steps that run automatically apart from setting the parameters and starting the module, are shown with black borders.

architecture. Note that while data acquisition and segmentation are critical steps in the processing pipeline, they are not part of our system.

4.3 Visual Computing Work-flow

The user work-flow is illustrated in Figure 4.6. After data acquisition and segmentation, which are not part of BRAINXPLORE, the user starts the pre-processing module where data is automatically converted to BRAINXPLORE’s internal data format.

The converted data can be inspected by the user and is then submitted to the index server. Index generation runs automatically. While the index is being built, the user performs the target definition and sets the sampling density for trajectory generation. These parameters are submitted to the trajectory pre-processing module where the risk signature of each potential trajectory is computed.

The user loads the resulting trajectory file into the user interface and defines rules for trajectory selection. Hard rules lead to the exclusion of trajectories that do not meet the requirements while soft (i.e., evaluation) rules are used for color-coding of candidate trajectories. After the hard rules have been defined and applied, the soft rules are evaluated for the remaining trajectories. Across different views, a consistent coloring scheme for trajectories is employed. In the Visual Analytics View, trajectories can be selected and filtered. Details on selected trajectories are presented to the user. The user

4. VISUAL COMPUTING WORK-FLOW AND SYSTEM DESIGN




Figure 4.7: Interaction of different views. Top left: Parallel Coordinates. Top right: risk signature chart used for trajectory inspection. Bottom left: the main 3D view Bottom middle: Supporting 3D View. Bottom right: Augmented Slice View.

inspects the set of remaining trajectories and refines the rules as needed. This iterative process is repeated until a feasible result has been reached. Furthermore, the trajectories themselves can be modified and evaluated interactively by querying the index server. In the end, a final decision is taken and a trajectory is chosen. Figure 4.7 illustrates the process. First, trajectories are selected in a PC view. Those trajectories are presented to the user in different 2D views, 3D views and visual analytics views, which will be introduced in Section 6.4.

Methods and Data

BRAINXPLORE is a system that consists of different modules which all fulfill specific roles. In the following sections we offer an overview of the methods we have employed in each module. The interaction of the modules with each other is illustrated in Figure 4.5.

The methods we used for data acquisition and segmentation are presented in Section 5.1. In Section 5.2, we outline the steps we have undertaken to preprocess the segmented data. Indexing is discussed in Section 5.3. Our methods of trajectory generation are presented in Section 5.4. Finally, in Section 5.5, filtering and refinement of generated trajectories is discussed.

5.1 Data Acquisition and Segmentation

We used a real dataset from a past biopsy to test and evaluate our system. The patient underwent contrast enhanced CT, MR and MRA scans. Also, a MRI scan was taken post-operatively to verify the site of the biopsy and exclude the existence of hemorrhage. All images exhibited different spatial resolutions. The image dimensions for each image modality is presented in Table 5.1.

Any segmentation software can be used to segment risk structures. In our case study, we used FreeSurfer for cortical and subcortical segmentation. Results for vessel segmentation

Data	Width	Height	Depth (Slices)
Pre-Op MRI	288	288	205
Post-Op MRI	232	256	192
MRA	528	528	125
CT	512	512	164

Table 5.1: Spatial resolution of the considered input data.

tend to be of insufficient quality for precise planning. Therefore, vessels are segmented separately in MeVisLab, from a combination of contrast-agent enhanced MRI and MRA data. The pyramidal tract was manually segmented by the neurosurgeon from MRI data. Finally, using MeVisLab, the tumor was segmented manually from the MRI and the skull was extracted from the CT. The acquired binary masks are then passed to the data preprocessing module. Although arbitrary data modalities can be included in the system, the minimal requirements of BRAINXPLORE are an MRI and a CT for the initial trajectory generation as discussed in Section 5.4.

5.2 Data Preprocessing

In contrast to other systems, we do not convert our data to polygonal meshes. Rather, we work directly on volume data and avoid additional computation time and conversion-related loss of precision.

The minimal distances to structures at risk defines a trajectory's risk signature. Those minimal distances are calculated based on the results of the data pre-processing step. Stray misclassified voxels can severely impair the end result. Hence, on each relevant structure, a connected component analysis is performed and only the largest connected component is considered for further processing. Each structure is then converted to a distance field, i.e., the distance to the closest non-zero voxel in the mask is computed for each voxel. Distances are capped to a pre-defined upper limit, as any distance further than that is considered safe for all practical purposes.

5.3 Indexing

As explained in Section 4.1.2, the neurosurgeon's work-flow demands quick access to spatial information. This is implemented via the use of a spatial index server as described by Bruckner et al. [BŠG⁺09]. The server stores spatial information for an arbitrary number of structures out-of-core and allows quick access to distance information. It is based upon two parts, i.e., a lookup volume, which is kept in memory and an out-of core distance table. Each voxel in the lookup volume stores an offset into the distance table and a number N of proximal items. The distance map stores, for each voxel P , a list of pairs of the form $\langle \text{distance}, \text{identifier} \rangle$, sorted by ascending distance to P . In order to enable efficient caching for subsequent accesses to the distance table, the entries of the distance table are arranged based on a three-dimensional Hilbert curve. Figure 5.1 illustrates the storage concept. When issuing a distance query, a simple lookup at the corresponding entry in the distance table yields a list of all proximal structures, together with their respective distances from the queried voxel.

The results of the indexing process are persisted to the hard drive for later use. Alternatively, the indexing server can also be interactively queried during runtime after a trajectory position has been refined in order to obtain updated distance information for




Figure 5.1: Index server storage concept as described by Bruckner et al. [BŠG⁺09]. To the left is a simple 2D scene containing three objects. In the middle, the lookup volume is illustrated. The light-grey overlay shows the corresponding Hilbert curve. To the right is the corresponding distance map.

the new trajectory. This allows the neurosurgeon to alter existing trajectories or, if no suitable candidate can be found, add new custom entry points.

5.4 Trajectory Generation and Preprocessing

In the **trajectory preprocessing module**, the set of potential entry points is calculated. As an input, the module takes the spatial index to work on, a sampling density and a target area definition in the form of either a single point or a list of points.

We assume that the head is in the center of the CT data set and define its center as the center of a spherical coordinate system. The diameter in each dimension is chosen in accordance to the respective data dimension, i.e., for a data set with data dimensions of 256^3 , the diameter is 256. With a user-defined number of latitudinal and longitudinal steps we sample the coordinate system and convert the spherical coordinates to Cartesian coordinates, as described in Algorithm 5.1.

Since the skull is not a sphere, the points created in this manner will not lie on the skin surface. Therefore, for each potential entry point, a ray to each specified target point on the tumor is cast. This ray is then traced in the CT volume until a voxel value above a specified threshold is hit. This way, the entry point is guaranteed to be on the skull surface. To prevent trajectories that would pass through the eyes, ears or mouth to be created, the number of crossed bone layers is counted, similar to the approach presented by Beyer et al. [BHWB07]. Trajectories that do not cross **exactly one** layer are discarded. Algorithm 5.2 described the process, Figure 5.2 illustrates an example. This correction process may lead to the generation of new entry points on the skull. Figure 5.3 shows how this happens. For an original entry point on the spherical

coordinate system, three rays are cast to their respective tumor target point. Since each point intersects the skull surface at a different voxel, three entry points are generated. This means that specifying multiple tumor target points may inherently lead to a denser skull sampling.

After the correction, for each calculated potential entry point, a ray is cast to each specified target point on the tumor and the minimal distances to each structure along that trajectory are queried. Good, noiseless distance fields (i.e., such where only relevant connected components above a certain size threshold are stored) are of crucial importance here. While not visually disruptive, noise in the form of misclassified voxels can lead to significantly altered minimal distances.

An initial list of candidate trajectories is created from every possible combination of entry sample points and tumor target points. These candidate trajectories are saved to a file for further use in the visualization module.

The trajectory preprocessing module can also be interactively queried at runtime to create a list of new entry sample points at a higher sampling density around a promising candidate. This allows the neurosurgeon to assess promising regions where one or more suitable entry points have already been found.

Algorithm 5.1: Creation of geographic coordinate system for sample generation.

After the loop, each sample point can be identified with the tuple (latitude, longitude)

Input: Volume Size s_{xyz} in x/y/z, number of latitudinal steps lat_{steps} , number of longitudinal steps lon_{steps}

Output: List c of potential entry points on Sphere

```

1 for  $x \leftarrow \{stepsize * n | n \in \mathbb{N}, n \in \{-180, 180\}\}$  do
2    $lat \leftarrow \frac{x * \pi}{180}$ ;
3   for  $y \leftarrow \{stepsize * n | n \in \mathbb{N}, n \in \{-180, 180\}\}$  do
4      $lon \leftarrow \frac{y * \pi}{180}$ ;
5      $c_{x,y,0} \leftarrow (s_x - 1) * \frac{\sin lon * \sin lat}{2.0} + \frac{s_x}{2.0}$ ;
6      $c_{x,y,1} \leftarrow (s_y - 1) * \frac{\sin lon \cos lat}{2.0} + \frac{s_y}{2.0}$ ;
7      $c_{x,y,2} \leftarrow (s_z - 1) * \frac{\cos lon}{2.0} + \frac{s_z}{2.0}$ ;
8   end
9 end
```

5.5 Trajectory Filtering and Refinement

After the initial set of potential trajectories has been generated on the index server as described in Section 5.4, it is subjected to optimization as well as incremental refinement and filtering. If more than one target point on the tumor is specified, optimization aims at choosing, for each entry point, the best target point such that the resulting trajectory candidate best satisfies the rules laid down by the neurosurgeon. The optimization process

Algorithm 5.2: Correction of entry-points based on CT data. Given a set of trajectories T described by all Bresenham points t along the trajectory, the algorithm finds the first point on the skull and returns the corrected (i.e., cropped) trajectory that now goes from the tumor target point to the skull surface.

Input: List of potential trajectories T based on result of Algorithm5.1, CT image data C

Output: List c of potential entry points on Sphere

```

1 for  $t \in T$  do
2    $cnt \leftarrow 0$   $bonecnt \leftarrow 0$ ;
3    $insidebone \leftarrow 0$ ;
4    $skullpoint \leftarrow \text{Nothing}$ ;
5    $goodidx \leftarrow -1$ ;
6   for  $pt \in t$ ,  $t$  is the set of Bresenham points in the current trajectory do
7     if  $C_{pt} > 0$  then
8       if  $insidebone == 0$  then
9          $insidebone \leftarrow 1$ ;
10         $bonecnt \leftarrow bonecnt + 1$ ;
11      end
12     else
13       if  $insidebone == 1$  then
14          $insidebone \leftarrow 0$ ;
15          $goodidx \leftarrow cnt$ ;
16       end
17     end
18      $cnt \leftarrow cnt + 1$ 
19   end
20   if  $bonecnt \neq 1$  or  $goodidx < 0$  then
21     return  $\text{Nothing}$ ;
22   else
23     return  $\{t_i \mid t \in T, i \leq goodidx\}$ 
24   end
25 end

```




Figure 5.2: Example of entry Point Correction as described in Algorithm 5.2. Three entry point candidates (A , B , C) have been sampled on the geographic coordinate system's upper hemisphere. Between point A and the tumor (yellow), exactly one layer of bone is passed. Hence, A is translated to the first intersection with the skull and stored as new entry point A' . Point B yields a trajectory that passes through two layers of bone. The trajectory between point C and the target point does not pass any bone layer. Hence, both entry points are discarded.

is described in Section 5.5.1. The resulting set of trajectory candidates is evaluated against hard constraints and filtered accordingly. The idea and process of this step is explained in Section 5.5.2.

5.5.1 Target Optimization

During the data loading process, the most suitable target point is selected for each potential entry point. In the GUI, up to three evaluation criteria can be specified for this optimization. Amongst these criteria are minimal distances to each structure, path length, insertion angles at the skull and tumor, etc. The criteria are processed hierarchically for pairs of trajectories. If two trajectories originate from the same entry point candidate on the skull but reach different target points on the tumor, the first criterion is evaluated for both trajectories. The results are then compared and a choice is made depending on a user-specified operation (i.e., $<$, $=<$, $>=$, $>$). If the trajectories are equal with regards to the first criterion, the second one is evaluated likewise. Finally, if both the




Figure 5.3: Side effects with the Entry Point Correction algorithm. A given entry point that is connected with a set of tumor target points may result in the generation of new entry points on the skull.

first and second criterion are equal for both trajectories, the third one is evaluated. If all three criteria are equal, the first trajectory supplied to the algorithm is chosen by default. Algorithm 5.3 describes the decision process. In Figure 5.4, an example is illustrated.

Algorithm 5.3: Optimization of candidate trajectories.

Input: Two candidate trajectories (T_1 and T_2) originating at the same entry point on the skull but reaching different target points on the tumor

Output: Better candidate trajectory based on evaluation criteria

```

1 for  $c \leftarrow \text{criterion}_{1,2,3}$ ,  $o \leftarrow \text{operation}_{1,2,3}$  do
2    $val_1 \leftarrow c(T_1)$ ;
3    $val_2 \leftarrow c(T_2)$ ;
4   if  $o(val_1, val_2)$  then
5     return  $T_1$ ;
6   end
7   if  $o(val_2, val_1)$  then
8     return  $T_2$ ;
9   end
10 end
11 return  $T_1$ 

```

5.5.2 Hard Rule Definition

The definition and application of hard rules aims at reducing the set of potential trajectories to a manageable number before evaluating their fitness based on **soft rules**. Trajectories that are longer than a user-specified length-threshold are discarded. Finally, the minimal distance to all structures along the trajectory is considered. Also, trajectories

that do not satisfy the user-defined safety requirements are discarded. Since all data has already been preprocessed by the trajectory generation module, distance look-ups can be performed in a very time-efficient manner. Algorithm 5.4 describes how the list of initial trajectory candidates is processed.

Algorithm 5.4: Pre-Filtering of target points.

Input: Unfiltered List ts of Trajectory Candidates, List of structures ss and corresponding minimal distances md

Output: Filtered List of Trajectory Candidates

```
1 for  $t \leftarrow ts$  do
2   for  $s \leftarrow ss$  do
3     if  $mindistance(s, t) < md_s$  then
4       discard  $t$ ;
5     end
6   end
7 end
8 return filtered List of Trajectory Candidates
```

5.5.3 Soft Rule Definition and Evaluation

In contrast to hard rules, which aim at reducing the number of potential trajectories by imposing rigid exclusion criteria, a set of soft rules can be used to evaluate the remaining trajectories and optimize the result. The quality Q_t of a trajectory t is defined as shown in Equation 5.1.

$$Q_t = \sum_{i=1}^n \frac{d_i * w_i}{max_i} \quad (5.1)$$

d_i denotes the minimal distance to structure i along the trajectory, w_i is the weight assigned to that structure and max_i is the maximal minimal distance found in all remaining trajectories. All weights need to sum up to 1.0. Each Q_t is then divided by the best encountered quality. This yields a relative measure, i.e., trajectory qualities are given with respect to the best currently found trajectory.




Figure 5.4: Example of target point optimization as described in Algorithm 5.3. Three trajectories originating from the same entry point but ending at different points on the tumor are sequentially compared. Optimization is set to distance maximization, with the distance to the nearest vessel being the top priority, followed by the distance to the nearest sulcus and finally, the distance to the ventricle. Sequentially evaluating the candidate trajectories against these optimization criteria yields the selection of Trajectory 3.

Implementation

To demonstrate the feasibility of our system, we have implemented a prototype. All modules were implemented and tested on a Windows 10 operated laptop PC with an Intel Core i7-3630QM processor, 12GB RAM and an NVidia GTX 660M graphics processor with 2GB dedicated GPU memory.

In this Chapter the implementation details of our system are presented. Section 6.1 briefly describes the MeVisLab work-flow. Specifics on data preprocessing are discussed in Section 6.2. Our implementation of the spatial index is presented in Section 6.3. The user interface, including details on each step of the user-work-flow is discussed in Section 6.4. Information Visualization techniques can be found in Section 6.5.

6.1 MeVisLab Work-Flow

MeVisLab is a rapid prototyping system developed by the MeVis Medical Solutions AG. It offers extensive data processing and visualization capabilities and is therefore an excellent choice for creating a prototype. The work-flow in MeVisLab is based on creating networks from various modules. Each module has *input* and *output* fields as well as *parameters*. Input and output fields can be image data, OpenInventor scenes, and base objects (e.g., lists of markers, strings, etc). Modules are placed with drag-and-drop and input fields are joined to the corresponding output fields of modules at an earlier stage in the network. Data evaluation is performed on demand, i.e., modules may or may not actually request data from their predecessors. Parameter fields communicate in the form of string variables. Figure 6.1 shows an example network in MeVisLab. New modules can be created by joining existing ones into macro modules and extending their functionality by adding Python [Ros95] or Java Script [Fla98] Code. Alternatively, new modules can be written in C++.

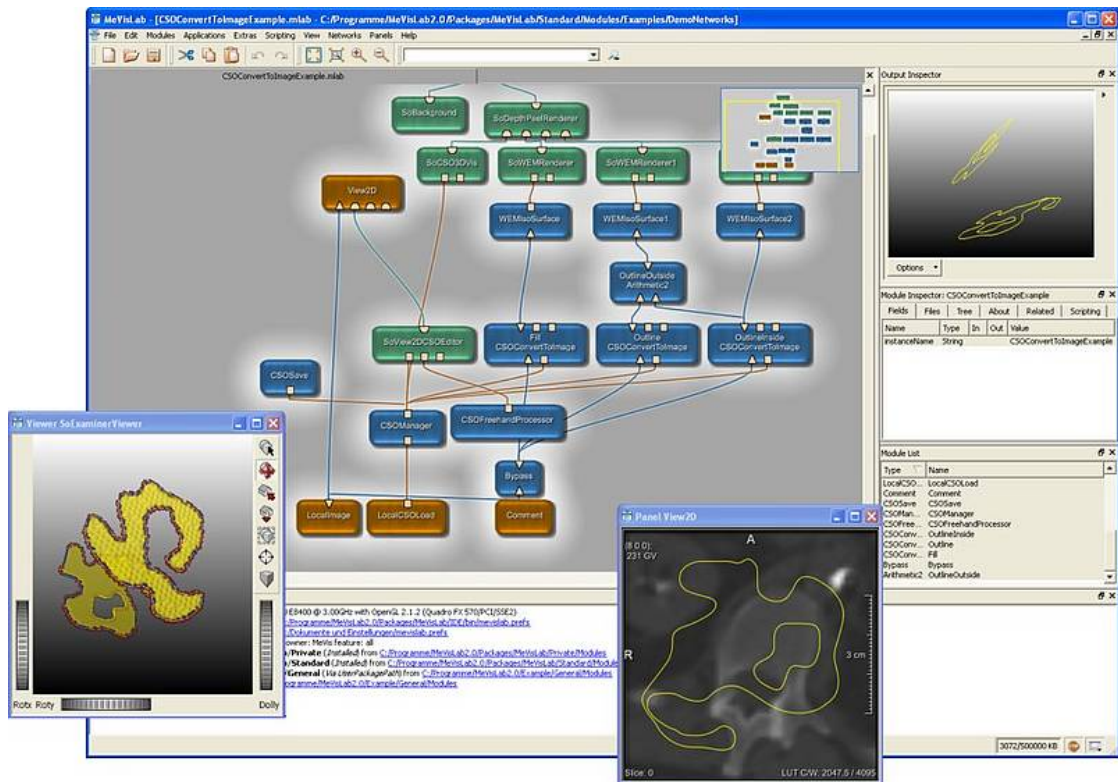


Figure 6.1: Example network in MeVisLab. Green nodes represent inventor modules, blue nodes are MevisLab modules and brown nodes are macro modules. Triangular connectors represent image data, round connectors represent Open Inventor Scenes and square connectors are used for Base objects such as lists. [AG17]

6.2 Data Preprocessing

Data preprocessing was performed in MeVisLab. Structures that were segmented with FreeSurfer and stored as NIFTI files were imported and processed. Connected component analysis was performed with MeVisLab's built in module *ComputeConnectedComponents*. From the resulting list, the largest component was filtered using the *FilterConnectedComponents* module.

Each structure was stored as a set of files. On the one hand, a DICOM file was stored for visualization purposes. On the other hand, a raw file in an 8 bit unsigned int format was stored for the corresponding distance field. Distance fields were calculated using MeVisLab's built in *EuclideanDTF* module. They were then capped to a maximum distance of 20 mm, as higher distances are not relevant for our purposes. Distances were scaled to the range [0,255] in order to fit the 8 bit data format.

6.3 Indexing

The spatial index server and the trajectory pre-processing module build upon existing work and are implemented in the Go Programming Language [DK15]. As pointed out by Hansen [Han11], "the data types may have a big influence on the size of the data file". In our prototype, distances are stored as unsigned 8 bit integer values and scaled to the range [0, 255] on the client side. Distances above 20 mm are not stored.

Initial trajectory generation can be executed directly from the user interface or alternatively, by usage of a command line program. In the former case, the results are directly communicated to the user interface by means of Python's *subprocess* package. In the latter case, a corresponding data file is stored in the CSV format for future use. In the file, each line represents a potential trajectory, identified by the entry point at the skull and the tumor end point. Additional columns indicate path length, insertion angle at the skull and insertion angle at the tumor. Finally, each line contains a list of structures at risk and their minimal distances in mm (or overlaps in voxels) from the trajectory.

The spatial index server has been implemented with a web interface and thus allows for a distributed work-flow. When queried, the server returns the risk signature of a single trajectory (i.e., a list of tuples where each tuple represents one structure at risk and the corresponding distance from or overlap with the trajectory) as a json object.

6.4 User Interface

The user interface was implemented in MeVisLab, based on the analysis presented in Section 4.1.2. Additional functionality was implemented with Python scripts. The system also uses OpenGL [SSK13] with GLSL vertex and fragment programs for visualizations that are not readily available in MeVisLab.

The **visualization module** displays trajectories and their contextual information with various views. The main view consists of a **Main 3D View**, a **Supporting 3D View**, an **Augmented Slice View**, a **trajectory query view** and an **Visual Analytics View**.

The Main 3D View is configurable and presents the trajectories in an anatomical context. Different data modalities (CT, MRI, the vascular system) are presented and can be toggled according to the surgeons needs. Candidate trajectories are represented as glyphs and color-coded according to their quality. Green markers denote the entry points of the highest quality, yellow markers indicate medium quality and orange markers indicate low - but still acceptable - quality. The weights for the cost function can be customized by the user. The currently selected trajectory is highlighted and the needle is shown as a blue line.

The Supporting 3D View displays the currently selected trajectory in a reduced context. Two modes have been implemented. In the first mode, the vascular system is color-coded with respect to the distance to the needle. Red areas are close to the needle

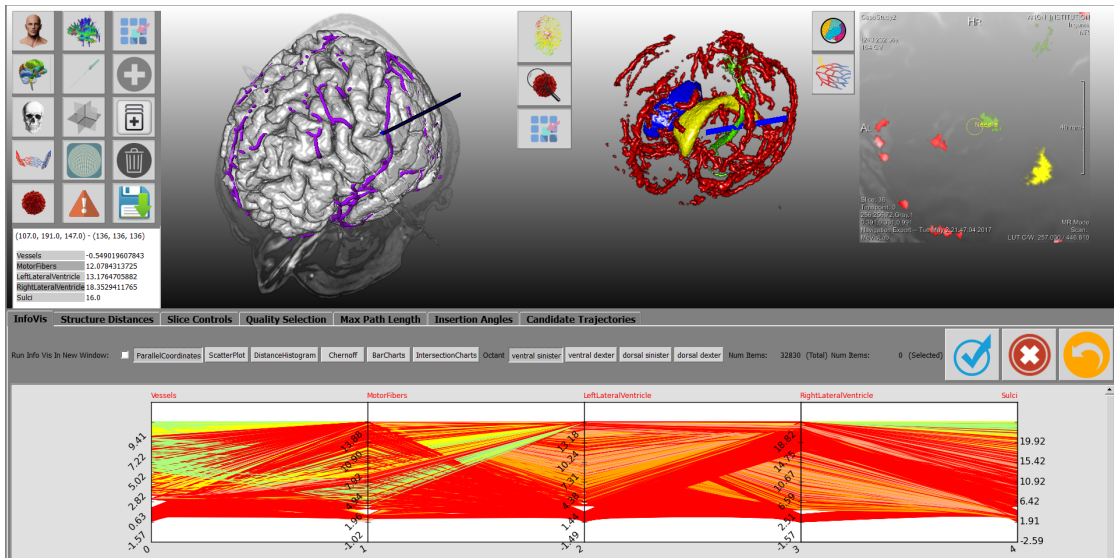


Figure 6.2: Visualization module user interface, with alternate supplemental views. Top left: Main 3D View. Top middle: Supporting 3D View, set to critical region view mode. Vessels are red, the left lateral ventricle is yellow, the right lateral ventricle is blue, and the pyramidal tract is green. Top right: Augmented Slice View, set to critical region view. Coloring is consistent with the Supporting 3D View. Bottom: Parallel Coordinates for trajectory picking. Between the Main 3D View and the command buttons, in the white box on the left, is the trajectory query Visual Analytics View that offers real-time information on the minimal distances for trajectories that have not been generated in the pre-processing step.

and therefore at risk, while green areas are further away and therefore safe. We used a MevisLab module for the distance calculation and visualization. In the second mode, selected structures at risk can be displayed. Each structure can be assigned a custom color and opacity value. Figures 4.7 and 6.2 illustrate the supporting 3D views.

The Augmented Slice View always presents a needle-aligned slice stack with additional annotations and highlights. In one mode, the augmented 2D view highlights vessels in the current slice and annotates the distance to the needle center. This is similar to the work performed by Herghelegiu et al. [HMP⁺12]. In another mode, the augmented 2D view highlights the structures at risk. Figures 6.2 and 4.7 illustrate the two modes of the augmented 2D view.

The Trajectory Query View offers a tabular overview of the best trajectory at a given point on the skull, similar to the work done by Bruckner et al. [BSG⁺09]. Information is queried online and can be obtained for arbitrary positions, even if no samples have been generated at this point during the initial sampling process. In Figure 6.2, the trajectory query view can be seen under the command buttons, in the white box on the left.




Figure 6.3: Manual Target Area Definition. Upper Left: 3D Contextual 3D view. The sagittal, axial, and coronal mid-planes are displayed. Highlighted in yellow is the segmented tumor. The red sphere indicates the target area's position and extent. Top Right, Bottom Left, Bottom Right: MPR view for the main anatomical planes, i.e., axial, sagittal, coronal.

The Visual Analytics View has been implemented using the Matplotlib library and operates in two distinct modes. In the global filtering mode, it can be used for trajectory selection. Three techniques have been implemented. These are discussed in Section 6.4.2. In the local evaluation mode, a trajectory's risk signature is presented. Two different views have been implemented and are discussed in Section 6.4.5.

Figure 4.7 offers an overview of the most important views in the user interface and their interaction during the planning process.

6.4.1 Target definition

After data segmentation and before trajectory pre-processing, the target area is defined in the user interface. The surgeon is presented with an MPR view of the volume. For better usability, the tumor is highlighted. As illustrated in Figure 6.3, the initial target location (represented by a red sphere in the left upper view) is selected in an orthogonal slice view and visualized in a contextual 3D view. A target radius in voxels can then be specified with a slider control. A sphere of the specified radius at the initial target point contains all final target voxels, i.e., all those voxels which are closer to the target than the specified radius. Each potential entry point is automatically connected with each target point and hence, the set of initial candidate trajectories is generated.




Figure 6.4: Trajectories are binned according to their minimal distance to configured structures. Bins of trajectories can be picked or rejected. Red bars indicate that a bin has been picked, while all trajectories in green bins will be rejected after filtering is applied.

6.4.2 Hard Rule Definition and Evaluation

Two methods of hard rule definition have been implemented. **Pre-filtering** is meant as a first, crude reduction step to exclude trajectories that violate minimal safety requirements. It is performed in the GUI by means of a set of slider controls which can be used to specify the the minimum distance to each structure at risk. After each filtering operation, the user is presented with the number of remaining trajectories that satisfy the rules. All operations can be undone in such a matter that the initial state (or any intermediate state resulting from filtering) can be restored at any point. **Post-filtering**, as opposed to pre-filtering, aims at the incremental refinement of rules and a step-wise approximation of the optimal result. The process is interactive in nature and offers different visual feedback mechanism. Filterable trajectory attributes are, among others, the minimal distance to each registered structure along the trajectory, path length, insertion angle at skull surface (deviation from a perpendicular insertion), insertion angle at tumor (deviation from a perpendicular insertion) and the distance between entry point and the cerebral cortex.

Different visualization and interaction methods to support post-filtering are provided. On the one hand, standard GUI elements such as sliders can be used to set scalar values (i.e., lengths, angles, distances to a single structure). On the other hand, InfoVis techniques are employed to provide insight into multivariate data (e.g., correlations between trajectory attributes are visualized in Parallel Coordinates, data distribution is visualized as a set of interactive histograms) and enable intuitive trajectory picking.

Four such picking techniques have been implemented. **Parallel Coordinates** are employed to present the user with all data dimensions in one view. Trajectories are represented as color-coded lines in the Parallel Coordinates. The coloring scheme is

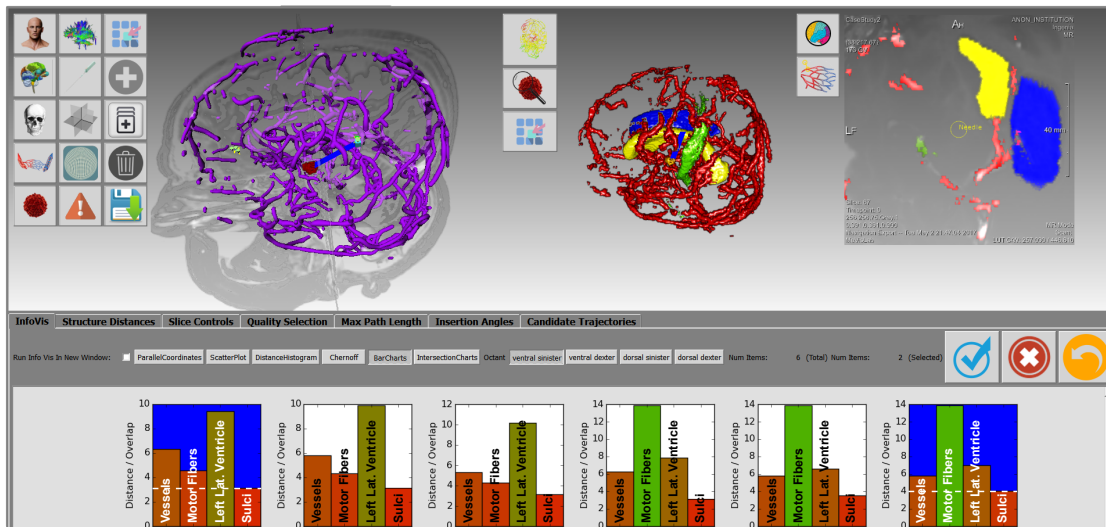


Figure 6.5: Multiple small bar charts illustrating each trajectory’s risk signature. Trajectories can be compared to each other and then accepted or rejected in this view. Colored bars indicate distances toward critical structures. Green bars indicate high distances, red bars indicate low distances and intermediate colors represent medium distances. The domain expert can rank trajectories based on his evaluation of the situation. Risk signature charts on blue background represent picked trajectories, those on white background will be rejected after filtering is applied.

consistent with the main 3D view. The Parallel Coordinates in the context of the user workflow is presented in Figure 4.7. **Configurable scatterplots** can be used for picking trajectories based on the most relevant data features. An arbitrary number of scatterplots can be displayed. **Distance histograms** allow the surgeon to pick one or more bins of trajectories for every data dimension. Each bin contains a certain set of trajectories. The union of all these sets is formed for each structure. Then, the intersection of all these sets is selected for further examination. This view can be used to prevent a filtering operation from returning an empty set of results. Figure 6.4 illustrates the concept. In the presented example, all trajectories with a distance of at least 4 mm to the closest vessel, 8 mm to the pyramidal tract, 9 mm to the left lateral ventricle and 7 mm to the closest sulcus have been selected, as indicated by the red color. In the main 3D view, selected trajectories are highlighted with light blue markers. This selection yields the same results as minimal distance thresholding, but allows for a better estimate of how many trajectories will remain after each filtering operation.

We also introduce novel ways of filtering trajectories based on their risk signature. **Multiple small bar charts** illustrating each trajectory’s risk signature are presented to the user, as illustrated in Figure 6.5. In each bar chart, color coded bars indicate the minimal distance to a structure at risk along the trajectory, with the highest bar indicating the largest distance and hence, the lowest risk. While this may seem counter-

intuitive at first, the reasoning was that the highest bar inherently draws the user's attention. Representing higher risks with higher bars emphasizes the trajectories that exhibit a greater risk. The goal, however, is to find the best suited trajectories for further evaluation, not identify those with the highest risk. Also, the color-coding, which goes from green for trajectories of low risk to red for those with a high risk further aids the selection process. Note that the charts are scaled with regards to the highest encountered distance and hence, scales may differ between individual charts. To illustrate the difference in scale, a white dashed line has been inserted at the 3 mm and 4 mm mark of the first and last trajectory.

Here, the color coding offers a more intuitive and concise measure than bar height alone. This becomes visible, for example, when comparing the second and fourth chart in Figure 6.5. The third bar, i.e., the distance to the left lateral ventricle, appears to be of the same height in both charts. Due to differences in scale, however, the distance in the second chart is 10 mm, while it is 14 mm in the fourth chart. This difference becomes evident when also considering the color coding, i.e., the bright green in the fourth chart as opposed to the olive green in the second chart. In Figure 6.5, the first trajectory maximizes the distance to the vessels. In contrast, the last trajectory does not perform as well with regards to vessels, but exhibits a higher safety margin towards the nearest sulcus.

An icon based view that maps trajectories to cartoon faces, as proposed by Chernoff [Che73] has been implemented, as shown in Figure 6.6. Trajectory properties such as length, distance to vessels, etc can be mapped to spatial features such as, e.g., the slant of the eyes or the length of the nose. Each of these views can be used to accept or reject trajectories based on their properties.

Potential trajectories are also visualized in a main 3D view. Figure 6.7 presents Parallel Coordinates used for selecting trajectories, along with the main 3D view. After each operation on the trajectory space, the set of options is refined accordingly and the number of remaining candidates is presented to the user. As in the pre-filtering step, every operation can be undone at any point.

6.4.3 Soft Rule Definition and Evaluation

After each filtering operation, all remaining trajectory candidates are evaluated for their fitness as described in Section 5.5.3. Rules can be defined in a configuration file where each line represents a soft constraint that will be used for quality evaluation. For each soft constraint, the corresponding weight is supplied.

Those trajectory candidates, which are not excluded due to the violation of a hard constraint are divided into six bins according to their quality. The best bin gets assigned the color green, the worst bin is red. The potential entry points in the Main 3D View are colored accordingly. Lines in the Parallel Coordinates are given the same colors as the entry points they represent. The coloring scheme can be seen in Figures 4.7 and 6.7.

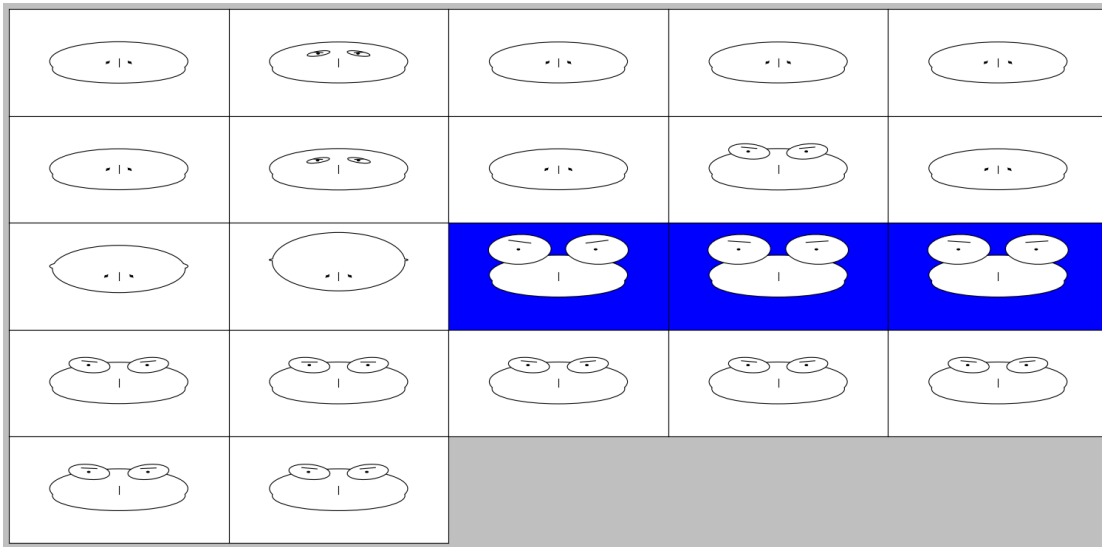


Figure 6.6: Chernoff faces map features of multivariate data to facial features on comic faces. According to Chernoff, by presenting data in a very familiar fashion, namely, human faces, "an investigator can quickly comprehend relevant information and then apply appropriate statistical analysis" [Che73]. Faces on blue background represent picked trajectories, those on white background will be rejected after filtering is applied.

6.4.4 Trajectory Addition and Removal

Trajectories can be added interactively at any time. This can be done in three ways: addition of new samples in the main 3D view, resampling around existing samples and finally, refinement of existing trajectories.

In the main 3D view, trajectories can be **added manually** at any time. A GPU-based risk evaluation was implemented in order to aid the surgeon in the decision where to place a new sample. Figure 6.8 presents the view. For each point on the skull surface, a ray is cast to the specified target point on the tumor. A distance field containing the vessel system is supplied to the GPU as a 3D texture. The texture is sampled at each point along the trajectory and the minimal distance to the closest vessel along the trajectory is found. The corresponding value is then projected onto the skull surface. Safe entry points are displayed in green, while entry points that result in the intersection of a vessel are displayed in red. This is an extension of the work done by Herghelegiu et al. [HMG11], who project the stability map of a given, user-defined ellipsoidal entry region onto a sagittal plane.

In Figure 6.8 A, a trajectory in a green, i.e., safe entry zone is presented. Conversely, Figure 6.8 B shows a trajectory in a red, i.e., unsafe entry zone. As can be verified in the supporting 3D view (top right) and Visual Analytics View (bottom), the minimal distance to the nearest vessel is much higher for the trajectory presented in Figure 6.8 A, than for the one presented in Figure 6.8 B.

6. IMPLEMENTATION




Figure 6.7: Trajectory evaluation after minimal hard rule definition. Trajectories that are closer to the closest vessel than 3 mm have been rejected. The weights used for evaluation are 0.5 for vessels and 0.25 for Sulci and the left lateral ventricle respectively. Trajectories are represented as glyphs in the main 3D view (top left) and consistently colored lines in the Parallel Coordinates (bottom).

Because of GPU limitations, only one tumor target point can be considered in our implementation. Furthermore, since GPU memory is limited, as explained in Section 3.1.5, only the vascular system is considered in this view.

New trajectories can also be **created on demand** by resampling around an existing, promising candidate trajectory. The trajectory preprocessing module is queried for a new set of trajectories. The newly created trajectories are then added to the set of currently available trajectories. Sampling works similarly to the initial sampling process described in Section 5.4.

Equation 6.1 describes the random translation of the spherical coordinates, where X is the random scalar used for translation, n_x, n_y, n_z are the corresponding coordinates of the newly created sample and o_x, o_y, o_z are the corresponding coordinates of the initial sample.

$$\begin{aligned}
 n_x &= X * radius * sign + o_x \\
 n_y &= X * radius * sign + o_y \\
 n_z &= X * radius * sign + o_z
 \end{aligned}
 \tag{6.1}$$

The coordinates obtained in this matter are adapted so that they are guaranteed to lie on

a bony surface as described in Algorithm 5.2 and converted into cartesian coordinates as shown in Algorithm 5.1. Figure 6.9 illustrates the process of adding new trajectories by resampling around an existing one. Steps that are performed in the visualization module are represented by black boxes, green boxes indicate a step is performed on the server module.

Finally, existing trajectories can also be refined in an MPR view. Both the entry point and the tumor target point can be translated. The spatial index server is queried for the risk signature of the new trajectory. A risk signature chart as illustrated in Figure 6.5 is presented to the user and the new trajectory is added to the set of currently available trajectories. This **manual refinement** function acknowledges the fact that, ultimately, the neurosurgeon is the domain expert and might make decisions based on criteria that are hard to quantify.

Exclusion of trajectories is usually done by refining the rules and applying the filter. All trajectories that do not satisfy the rules are excluded from further examination. However, if a trajectory does satisfy all rules but is still deemed unsuitable by the neurosurgeon, it can also be **deleted manually**.

6.4.5 Trajectory Inspection

During evaluation, the reduced set of trajectories is presented to the user. As in the post-filtering stage, the user can navigate between different views. The Main and Supporting 3D Views offer spatial context, while the Visual Analytics View presents detail information. A trajectory's risk signature is displayed in one of three ways in the Visual Analytics View.

The **Visual Analytics View** offers detail information on the trajectories risk signature. Figure 4.7, on the top right, illustrates the risk signature chart view. Bars are colored according to their value with a simple linear interpolation in RGB space between green (for distances of 20 mm) and red (for distances of 0 mm). In case a trajectory does intersect a structure, a negative bar representing the total overlap volume between the needle and the structure is presented. These bars are displayed in red.

Alternatively, the distance profile from the needle to all structures at risk can be plotted as a set of colored functions in a shared coordinate system. The spatial index server can be queried for all voxels that lie on the trajectory. Sample positions are calculated with a modified Bresenham's line drawing algorithm. Figure 6.10 illustrates the needle pathway distance graph view. Although the most important measure regarding the distance to a structure along the trajectory is the minimum, considering the distance at multiple sample points might reveal additional insight into the trajectory's quality. For example, although two trajectories might exhibit an identical minimal distance to a given structure, with everything else being equal, the trajectory with a higher average distance is better than the one with a lower average distance.

Finally, the radar chart view presents a trajectory's risk signature as a geometric figure. Radar charts can be freely configured to include arbitrary structures such as to best

suit the user's needs. On each plot, the trajectory is presented as a pair of shapes. One shape, the distance polygon, reflects the minimal distances to structures. The second shape, the overlap polygon, depicts the overlap with structures. Shapes are created by setting up a polygonal coordinate system where each data dimension is represented by an axis connecting the plot center with the respective polygon vertex. For example, a three-dimensional dataset can be represented in a triangular coordinate system, while a four-dimensional dataset requires a square or diamond-shaped coordinate system. Data points in each dimension are scaled such that each axis has the same length. In our implementation, the area under the distance polygon gives an estimate of the safety margin. Hence, distance polygons with larger areas and overlap polygons with smaller areas generally represent safer trajectories. In order to alleviate the radar chart's shortcomings described by Staffoglia et al. [SLF⁺11], namely, the inability of the radar chart to differentiate between the case when information on a data dimension is missing and the data dimension's value is zero, we fill missing data dimensions with a value of 20.

A distance of zero or less indicates an intersection with, and hence, a violation of a structure at risk. Representing missing values as zero therefore leads to the potential exclusion of trajectories. The only way a data dimension is missing, in our case, is if it is capped by the spatial index server. This happens if a structure is further away from the needle center than 20 mm. In that case, it is deemed not to be at risk, i.e., it is safe. Hence, while substituting 20 mm for the true value skews the data, this is of no clinical relevance.

Spatial information regarding the trajectories relation to structures at risk is also presented to the user. In a Supporting 3D View, the vascular system can be color coded with regards to the vessel's distance to the needle. Alternatively, the needle is presented in the context of selected critical structures, as shown in Figure 6.2. An Augmented Slice View presents the needle-aligned slice stack with highlighted vessels and distance annotations. Alternatively, the surgeon can choose to display a selected set of critical structures highlighted in the needle-aligned slice stack, as presented in Figure 6.2.

Finally, the results can be viewed slice by slice in either the original coordinate system or in a needle-aligned slice stack. These views are very similar to those offered by the system our domain expert currently uses. By creating familiar views, we aim at making the transition from one system to the other easy.

6.4.6 Trajectory refinement

Trajectories can be refined in two ways, either by **resampling** around an existing entry point or by **manually altering** the trajectory start- and endpoint. On the one hand, the trajectory generation module can be queried for a set of trajectories sampled at a higher resolution around an existing entry point as described in Section 6.4.4. This approach is meant to support the automatic addition of new trajectories in promising entry areas, at earlier stages of the planning process. Alternatively, the surgeon can interactively modify a trajectory in either the original coordinate system or in a needle-aligned slice stack

and query the spatial index server for updated distance information. The new trajectory is then added to the visualization module. This approach is meant to be used in later stages of the planning process, when acceptable solution may have already been found but expert knowledge makes certain subtle changes preferable.

6.5 Information Visualization

Although Staffoglia et al. [SLF⁺11] deem the spie chart the best representation for multivariate data for clinical outcome indicators, we have decided to rather implement the radar chart to illustrate the quality of individual trajectories.

This is due to the nature of our data. Rather than clinical outcome data, which is limited to a positive range, we have to display spatial information that can be either positive or negative, as described in Section 6.4.5. Also, the ambiguity between missing data dimensions and those items whose value is zero is a non-issue in our case, as missing information can be interpreted and handled easily in our system. The details of our radar chart implementation are discussed in Section 6.4.5.

Needle pathway distance graphs are implemented such that each structure at risk is represented by an individual line. This allows for a more detailed qualitative analysis of trajectories. Different structures may have different levels of risk associated with them, so displaying only a single line representing the nearest structure at risk for every sample point might take away important information. Figure 6.10 illustrates our implementation.

We have implemented Parallel Coordinates in which trajectories can be interactively picked and filtered. To the best of our knowledge, no existing solution in the context of brain biopsy planning implements this interaction technique.

Finally, Chernoff Faces seem to be a feasible visualization technique, as brain regions can be mapped to facial features based on their function. For example, the distance to the optic chiasm can be mapped to the size of the eyes, while Broca's area, the motor center for speech, can be mapped to the slant of the mouth. While not pre-attentive, in the scenario of brain surgery, this view enables to user to directly map at least a few features to the cartoon face. Figure 6.6 shows our implementation.

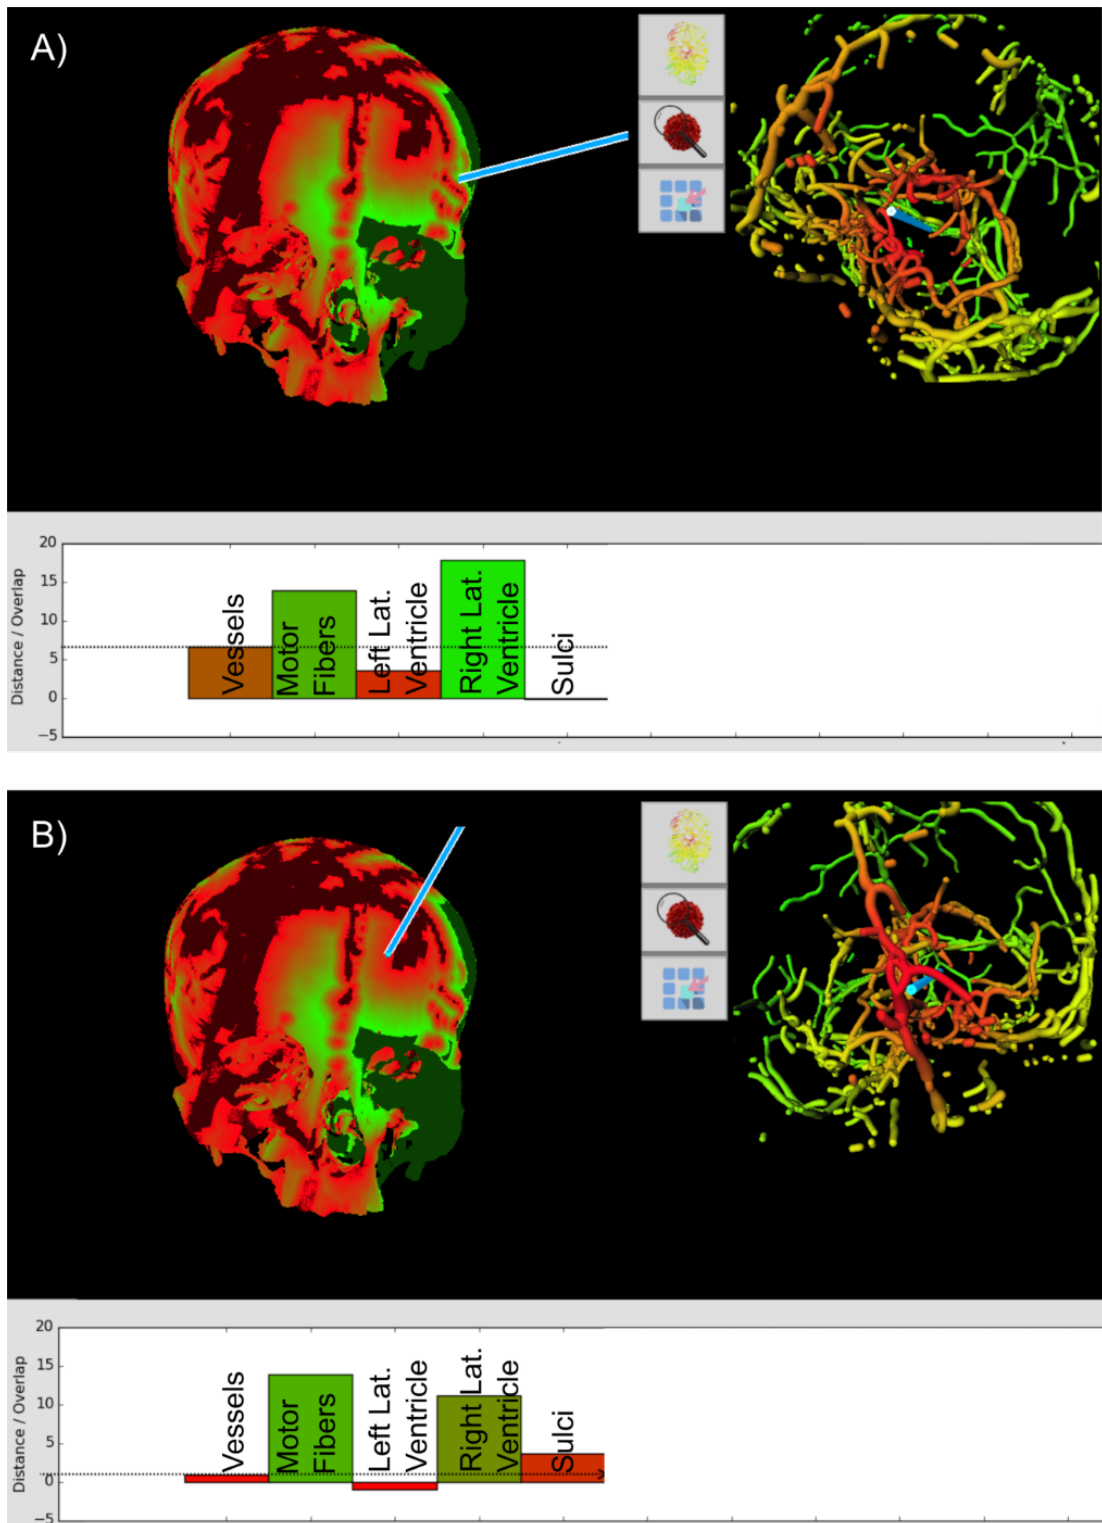


Figure 6.8: Presented in the Main 3D View (left) is the risk map. For each point on the skull surface, a ray is cast to the specified target point on the tumor and the minimal distance to the nearest vessel along the trajectory is found. Safe entry points with regards to the vascular system are displayed in green, while entry points that result in the intersection of a vessel are displayed in red. The Supporting 3D View (right) shows a color coded representation of the trajectory in the context of the vascular system. Blood vessels in red are closer to the trajectory, i.e., at risk. in the Visual Analytics View (bottom), the risk structure of the selected trajectory candidate is presented. A) Trajectory in a green, i.e., safe entry zone. B) Trajectory in a red, i.e., unsafe entry zone.

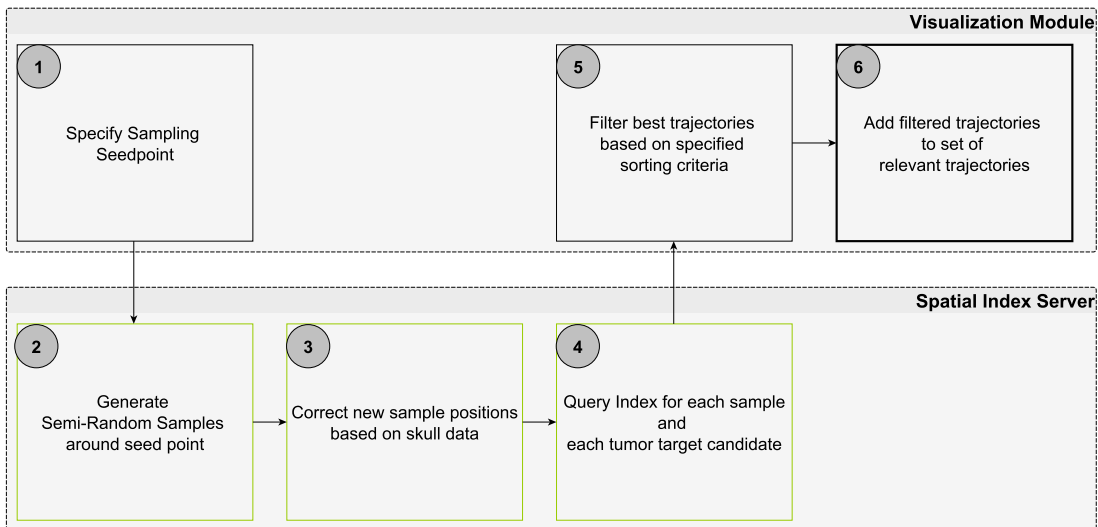


Figure 6.9: Adding new trajectories by resampling around an existing one. Black boxes: steps in the visualization module. Green boxes: steps that are performed on the spatial index server.




Figure 6.10: The needle pathway distance graph (bottom) represents the currently inspected trajectory (visible in the contextual 3D view in the top left and center) as a set of line graphs. Each structure at risk is represented by a colored line. The abscissa represents the current insertion depth while the ordinate shows the minimal distance to the respective structure at that position. This representation summarizes the information that could be gathered by slicing through the annotated volume in the augmented slice view (top right).

Evaluation

We performed an internal evaluation of BRAINXPLORE, where we evaluated the system for performance, accuracy and result quality. The methods used for the evaluation as well as the results are presented in Section 7.1. Furthermore, we performed an external evaluation and compared BRAINXPLORE to existing solutions. Our findings are presented in Section 7.2.

7.1 Internal Evaluation

We evaluated the usefulness of our application based on the dataset used for the case study in Section 4.1.1. Because our input data came from independent scans, pre-processing had to be performed. The CT and MRA scans were resampled in order to produce results with the same resolution as the FreeSurfer results (256^3). The resampled results were manually registered to the pre-operation MRI scan. Then, semi-automatic vessel and tumor segmentation was performed in MeVisLab.

Based on this data, different spatial index files were built. Candidate entry points were sampled on a spherical coordinate system at 255 latitudinal steps and 255 longitudinal steps, as described in Section 5.4. For trajectory generation, we provided the system with different tumor point configurations. The most basic configuration contained only one target point at the tumor's periphery. To test the performance for multi-target scenarios, we added additional definitions for 3, 5, 10 and 15 target points.

With those spatial index files, entry points, and target points, we evaluated the performance of the individual components of our system. Then, in cooperation with our domain expert, we used BRAINXPLORE to produce a biopsy trajectory for the dataset in our case study and evaluated the quality of the result.

The rest of this Chapter is structured as follows: First, in Section 7.1.1, we present the performance of our system under different conditions. Then, in Section 7.1.7, we show the neurosurgical evaluation of our results.

7.1.1 Performance Evaluation

We examined the efficiency of each individual part of our solution. A laptop computer running on a Windows 10 Professional 64 bit operating system with an Intel Core i7-3630 CPU was used for the evaluation. The machine has 12 GB of memory and an NVIDIA GeForce GTX 660M with 2 GB of video memory. In a real world scenario we expect the number of relevant trajectory candidates to be in the vicinity of 200-6000, as stated by Zombori et al. [ZRN⁺14]. However, use cases of upwards of 10000 trajectory candidates are possible. We have therefore decided to test our system with higher numbers of trajectory candidates than usually encountered in real world examples.

7.1.2 Index Generation

Although it is not within the scope of this thesis, the spatial index is a crucial component of our system. In order to assess the usefulness of our system, we have decided to include an evaluation for this component as well. To evaluate the performance of the spatial index generation, we measured the average computation time for the creation of indices of different sizes. Indices were built for 1, 5, 10, 50, 100, 150 and 178 structures at risk. The maximum number of 178 structures resulted from the extraction of cortical and subcortical regions of interest from our data set. We used the Destrieux Atlas provided by FreeSurfer for cortical parcellation the atlas developed by Fischl et al. [FSB⁺02], which is also provided by FreeSurfer, for subcortical segmentation. All empty results, i.e., those structures that could not be classified, were discarded. This process resulted in 175 structures. In addition to those structures, we used our semi-automatically segmented blood-vessels, a union of all sulci (except for the precentral and postcentral sulcus) as well as the manually segmented pyramid tract.

Each configuration was built five times and the times were recorded using the *Measure-Command* function of Microsoft's PowerShell. Table 7.1 provides details on the measurements, Figure 7.1 shows the average runtimes as a function of the number of structures at risk. As can be seen, the time required to construct an index for a single structure is approximately 13 seconds. This time is used for startup of the system and the allocation of memory.

7.1.3 Online Interactive Trajectory Query

To evaluate the performance of the spatial index query, we measured the average computation time for sending queries to the index server and processing the results on the client side. This excludes visualization. For the evaluation, different sets of tumor target points have been considered. We performed evaluations for 1, 3, 5, 10 and 15 distinct target points on the tumor surface. Additionally, we performed all tests on different spatial

Structures at Risk	1	5	10	50
Measurement 1	12440	13611	14931	25017
Measurement 2	12305	14058	15473	25122
Measurement 3	13247	13575	14732	25402
Measurement 4	13107	13847	15464	25284
Measurement 5	12647	13862	15048	25537
Average Runtime	12749	13791	15130	25272
Structures at Risk	100	150	178	
Measurement 1	48423	71189	94306	
Measurement 2	48319	72098	83914	
Measurement 3	47876	70832	85208	
Measurement 4	48742	71526	85948	
Measurement 5	48027	71208	84278	
Average Runtime	48277	71371	86731	

Table 7.1: Time measurements for index generation, in ms, with respect to the number of structures at risk.

indices, which were built on different numbers (1, 50, 100, 150, and 178, respectively) of structures at risk.

Each configuration was profiled by choosing 150 random points on the skull surface and sending a trajectory query to the server. For multi-target scenarios (i.e., all cases except the first), target optimizations were performed as described in Algorithm 5.3. We measured computation times and recorded the average for each configuration. Table 7.2 lists the resulting timings with regards to the number of structures at risk and the number of tumor target points. We included average runtimes across the different spatial indices to illustrate an average scenario.

Figures 7.2 and 7.3 offer a graphical representation of our findings. From the data we conclude that the runtime is linearly dependent on the number of considered tumor targets but not on the number of structures at risk. Due to the nature of the spatial index, access times are dependent on the sparsity of the data rather than the number of structures, because the trajectory is evaluated for each rasterized point. Passing through a higher number of voxels that have distance information associated with them leads to a higher number of calculations. The number of considered structures at risk, however, is not relevant to the computation time.

For real-time framerates, computation times need to be below $\frac{1}{30}$ of second, while interactive framerates call for computation times below $\frac{1}{20}$ of a second. From the data in Table 7.2, we conclude that on typical data sets with a resolution of 256^3 , our system can provide near real-time trajectory information for up to 10 targets on the tumor. Interactive framerates can still be achieved for up to 15 targets.

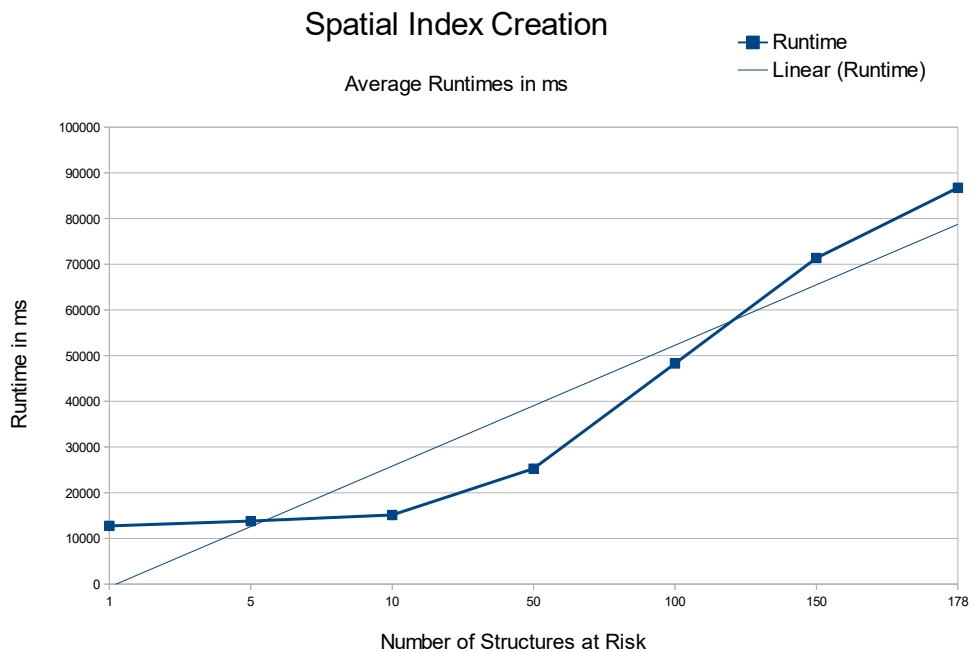


Figure 7.1: Average runtimes for the spatial index generation with respect to the number of considered structures at risk

Targets / Structures at Risk	1	50	100	150	178	Average
1 Target	4	4	5	5	5	5
3 Targets	7	8	8	8	8	8
5 Targets	12	12	12	12	13	12
10 Targets	23	23	24	24	25	24
15 Targets	33	34	36	35	36	35

Table 7.2: Average runtimes in ms for a trajectory query and subsequent optimizations. Also provided are the average runtimes across different numbers of structures.

7.1.4 Sample Point Generation

We also evaluated the performance of our system by profiling the sample generation process with different sets of tumor target points, as described in Section 7.1.3. We set up a spatial index consisting of 178 structures at risk. Then we generated a set of samples on a spherical coordinate system consisting of 250 longitudinal and 250 latitudinal steps as described in Algorithm 5.1. Each of created grid points was corrected as described in Algorithm 5.2. From each corrected entry point, a trajectory to each specified tumor point




Figure 7.2: Average runtimes for querying a trajectory with respect to the number of considered structures at risk and tumor target points. It can be seen that the runtime strongly correlates with the number of tumor target points.

was calculated on the server. First, we created trajectories for a single target point on the tumor. Then, we created sets of trajectories for 3, 5, 10 and 15 targets, as described in Section 7.1.3. For reasons explained in Section 5.4, introducing additional target points yields additional trajectories. Each target configuration was run 50 times and the average time was recorded. Table 7.3 offers an overview over the average computation times. Figure 7.4 illustrates the linear relationship between the computation time and the number of target points on the tumor.

The resulting set of trajectories was then parsed on the client and in the multi-target case, trajectory optimizations were performed as described in Algorithm 5.3. Also, trajectory length was calculated on the client. We evaluated the client with the same method as outlined above. Table 7.4 offers an overview of the average computation times. Figure 7.5 illustrates the results.




Figure 7.3: Average runtimes for querying a trajectory with respect to the number of considered structures at risk and tumor target points. It can be seen that the runtime strongly correlates with the number of tumor target points, but not with the number of structures at risk.

Targets	Trajectories	Sample Generation Time
1	19766	22
3	46468	59
5	62860	97
10	82158	192
15	89376	291

Table 7.3: Average computation times in ms for the initial sample point generation with respect to the number of tumor target points.

7.1.5 Soft Rule Evaluation

To evaluate quality evaluation times, an initial sampling resulting in 32830 trajectories was performed. This set was incrementally filtered by adding rules until only four trajectories remained. Each set was evaluated five times and the average runtime was recorded. The evaluation was based on four criteria, each of which contributed 25% to

Targets	Trajectories	Trajectory Optimization Time
1	19766	15
3	46468	49
5	62860	93
10	82158	176
15	89376	335

Table 7.4: Average computation times in ms for the initial sample point optimization with respect to the number of tumor target points.




Figure 7.4: Average runtimes for the initial sample point creation with respect to the number of tumor target points.

the overall trajectory quality score: the distance to the nearest vessel, the distance to the nearest sulcus, the distance to the left lateral ventricle and the path length. Equation 7.1 describes how the trajectory quality was calculated for a given trajectory i . D_j represents the trajectories distance to structure j , Max_j is the maximum distance to that structure amongst all considered trajectories. Trajectories are binned into quality categories for visualization purposes, as described in Section 6.4.3.

The results are presented in Table 7.5. Figure 7.6 illustrate the data. To fit all data, the

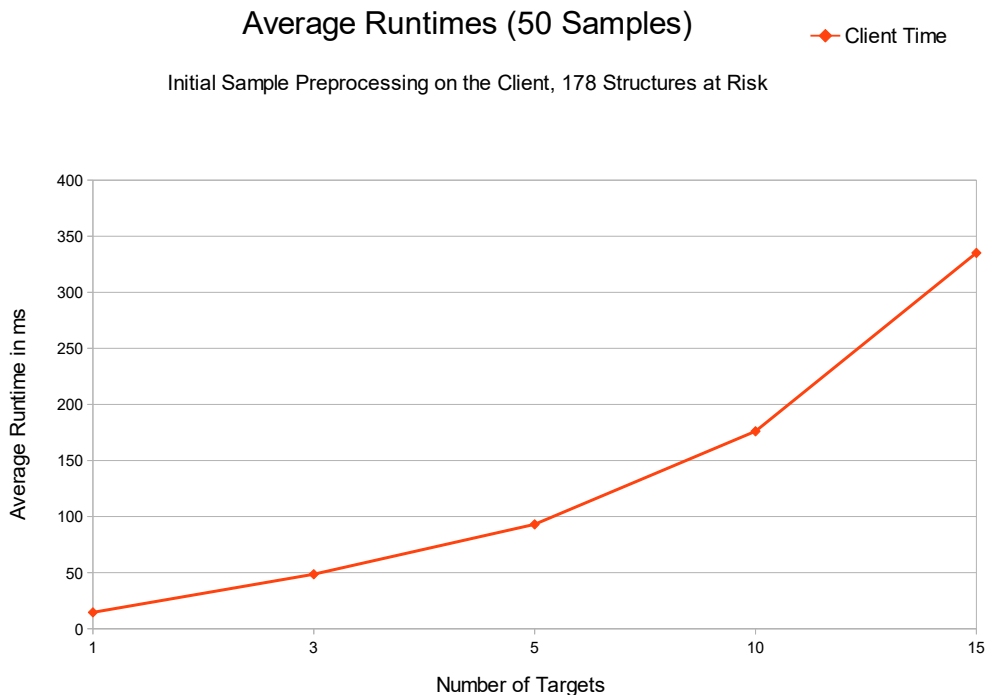


Figure 7.5: Average runtimes for the initial sample point optimization with respect to the number of tumor target points.

X-Axis was scaled logarithmically. As can be seen, the runtime is a linear function of the number of trajectories.

$$Q_i = \sum_{j=0,1,2,3} \frac{1}{4} D_j / Max_j \quad (7.1)$$

7.1.6 Rule Refinement

To evaluate the performance of our rule refinement algorithm, we tested it with respect to the number of rules to be evaluated. As described in Algorithm 5.4, our filtering algorithm discards any trajectory that does not satisfy all defined rules. We sped up the process by implementing early termination, i.e., the evaluation for a given trajectory candidate is stopped as soon as the first rule violation is detected. Therefore, our hypothesis was that performance is dependent on the order in which the rules are applied. To test our hypothesis, we created an initial sample space of 32830 trajectories and a set of 17 rules. For the first test, we applied the rules in a fixed order such that the most significant rules (i.e., those that lead to the highest number of discarded trajectories) are evaluated

Trajectories	Measure 1	Measure 2	Measure 3	Measure 4	Measure 5	Avg. Runtime
32830	1896	1890	1876	1873	1886	1884.2
17339	1173	1190	1180	1190	1167	1180
6081	582	586	586	587	588	585.8
3053	320	312	314	314	313	314.6
1812	200	202	199	203	199	200.6
588	122	117	117	118	116	118
106	30	25	25	26	26	26.4
1	3	3	3	3	3	3

Table 7.5: Average runtimes in ms for soft rule evaluation on a set of four rules. Each evaluation was performed five times, Measures 1 through 5 show the respective timings in ms.




Figure 7.6: Average runtimes in ms for soft rule evaluation on logarithmic scale.

in the beginning. We filtered our initial sample space iteratively and added one rule in each iteration, i.e., the same sample space was first filtered with a single rule, then with two rules and so on, until all rules were enabled. To test the stability of our approach, we ran the process five times and recorded runtimes for each individual test run as well as average runtimes. Table 7.6 presents the resulting times.

In Figure 7.7 it can be seen that the runtimes appear to be a linear function of the number of evaluated rules, until rule six is evaluated, at which time they exhibit a considerable drop-off. At this point, the combination of three rules (namely, distance to vessels, distance to the sulci and distance to the left lateral ventricle) lead to an early termination of the evaluation for all but 20 samples. This supports our hypothesis that the order of rules significantly affects the performance.

We performed the same experiment again, on the same set of initial samples. The second time, however, the order of rule application was randomized, i.e., in each of the five test runs, the order of rule application was different. Table 7.7 presents the resulting runtimes. As can be seen in Figure 7.8, the same characteristic drop-off in computation times can be observed in all measurements, only at different times. This indicates that the introduction of the respective rule led to a set of rules which result in the early rejection of most trajectories. We conclude that while the computation time is a linear function of the number of rules to be evaluated, the order in which those rules are introduced can be very significant.

7.1.7 Neurosurgical Evaluations

Our domain expert employs biopsy needles with a diameter of 3.5 mm and therefore a radius of 1.75 mm. We compute minimal distances from the needle center to structures at risk. The safety margin then describes the distance between the needle's edge and the structure, i.e., we subtract the needle radius from the minimal distance.

Pre-filtering was applied with the minimal distances presented in Table 7.8. This yielded 390 potential trajectories. Limiting path length to 80 mm reduced this set to 20 potential trajectories. Samples were further spatially limited to the left side, behind the hairline, by defining an appropriate ROI in the user interface. This reduced the set to 9 valid candidates. These were evaluated based on their minimal distance to the sulci.

From the set, the best trajectory (i.e., the one with the largest safety margin towards the sulci) is illustrated in Figures 7.10 and 7.11. Figure 7.10 presents the trajectory in an orthogonal slice view in the original coordinate system. The solid yellow line indicates the part of the trajectory behind the current needle tip position, dashed lines indicate the part of the trajectory that has not yet been crossed. As can be seen in the coronal view, the trajectory does not originate at a sulcus (highlighted in orange). The sagittal view reveals that the trajectory laterally passes the ventricle. Table 7.9 lists the minimal distances to structures at risk along the needle trajectory for those structures that are within the critical range of 20 mm. The associated risk signature is illustrated in Figure 7.9.

Rules	Measure 1	Measure 2	Measure 3	Measure 4	Measure 5	Average Runtime
1	1184	1160	1216	1215	1221	1199.2
2	1816	1851	1859	1864	1888	1855.6
3	2388	2500	2457	2485	2477	2461.4
4	2955	3003	3004	3023	3039	3004.8
5	3480	3590	3578	3603	3620	3574.2
6	1055	1085	1081	1090	1089	1080
7	891	914	906	910	899	904
8	875	891	899	891	913	893.8
9	871	897	887	893	898	889.2
10	906	917	916	906	926	914.2
11	912	944	944	939	942	936.2
12	922	960	943	938	942	941
13	944	970	958	952	943	953.4
14	955	970	962	950	957	958.8
15	965	1006	1011	992	995	993.8
16	994	1023	1013	999	1004	1006.6
17	985	1034	1021	1000	1024	1012.8

Table 7.6: Average runtimes in ms for rule refinement with respect to the number of evaluated rules. Each evaluation was performed five times, Measures 1 through 5 show the respective timings in ms. The rule order was not randomized.

Figure 7.11 presents the same trajectory, this time in an altered coordinate system obtained by performing multiplanar reformation along the trajectory path on the volume. For better spatial orientation, the postcentral sulcus is highlighted in a very light pink. This view presents the situation even more clearly - no sulci or vessels are crossed, high distances to the ventricle and pyramidal tract are maintained over the whole trajectory. Hence, all demands our domain expert places upon a suitable trajectory are satisfied.

We have reconstructed the reference trajectory from the post-operation MRI and evaluated it using our system. Figure 7.12 presents the trajectory in an orthogonal slice view in the original coordinate system. Table 7.10 lists the minimal distances to structures at risk along the needle trajectory. The corresponding risk signature is presented in Figure 7.9, at the bottom. As can be seen, the reference trajectory exhibits a higher distance to the precentral sulcus than the solution computed by our system. On the other side, distances towards the vessels, left lateral ventricle and all other cortical sulci are slightly higher in our solution.

Our domain expert has given us the feedback that BRAINXPLORE is a feasible planning tool, compared to manual planning on the Medtronic workstation and EpiNav [ZRN⁺14]. EpiNav is a software that was developed for pre-surgical trajectory planning for electrode placement in the context of deep brain stimulation and does therefore not work well for brain biopsy planning. With manual planning, under the assumption that all medical

Rules	Measure 1	Measure 2	Measure 3	Measure 4	Measure 5	Average Runtime
1	1172	799	1171	1181	1179	1100.4
2	1795	896	1803	1804	1798	1619.2
3	961	930	2409	2605	2002	1781.4
4	994	1022	3006	2654	2322	1999.6
5	950	1023	3532	2801	2828	2226.8
6	1022	1020	3742	2827	1051	1932.4
7	1032	1078	4381	3204	1102	2159.4
8	1101	1041	5077	3459	1129	2361.4
9	927	1091	2216	3739	1207	1836
10	922	1127	1967	1096	1283	1279
11	972	1141	2223	953	1279	1313.6
12	917	1145	2506	948	1355	1374.2
13	929	997	2585	995	1372	1375.6
14	941	986	2638	997	1268	1366
15	950	1011	2941	1006	1016	1384.8
16	944	1003	1046	1007	1009	1001.8
17	985	1024	1041	1016	1013	1015.8

Table 7.7: Average runtimes in ms for rule refinement with respect to the number of evaluated rules. Each evaluation was performed five times, Measures 1 through 5 show the respective timings in ms. The rule order was randomized.

Structure	Minimal Distance	Safety Margin
Vessels	4	2.25
Pyramid Tract	4	2.25
Left Lat Ventricle	4	2.25
Right Lat Ventricle	4	2.25
Sulci	3	1.25
Gyrus sup. temp.	4	2.25
Occipital Pole (LH)	4	2.25
Pars Opercularis (LH)	4	2.25
Pars Triangularis (LH)	3	1.25
Gyrus Supramarginalis	4	2.25
Postcentral Sulcus (LH)	4	2.25
Precentral Sulcus (LH)	4	2.25

Table 7.8: Minimal distance in mm to each structure along the needle trajectory in the case study. The safety margin is calculated by




Figure 7.7: Average runtimes in ms for rule refinement with respect to the number of evaluated rules. Each evaluation was performed five times, Measures 1 through 5 show the respective timings in ms. Since the rule evaluation order is fixed, early termination leads to significant drop-offs in computation time after the sixth rule is evaluated. At this point, only 20 samples are subjected to the full rule set.

imaging data is present on the Medtronic workstation, the planning process takes approximately 5 to 15 minutes. However, this is the time it takes to evaluate only a few possible trajectories and pick an acceptable one. BRAINXPLORE can calculate and evaluate a big number of trajectories in this time and hence, lead to potentially better results. Also, for novice users that do not have a high degree of experience, the planning procedure takes longer, approximately 15 - 30 minutes. For these users, BRAINXPLORE can on the one hand bring a considerable speed benefit. On the other hand, novice users can verify and evaluate their manual planning result by comparing it to the trajectories computed by BRAINXPLORE.

7.1.8 User Interface Evaluation

The color-coded glyphs have been particularly well received by our expert. Also, the trajectory inspection views illustrated in Figures 7.10 and 7.11, have proven useful. Contrary to our expectations, the Chernoff face view was not well received. We believe




Figure 7.8: Average runtimes in ms for rule refinement with respect to the number of evaluated rules. Since the rule evaluation order is randomized, early termination leads to significant drop-offs in computation time at different points. After the most significant rules have been implemented, only 20 samples are subjected to the full rule set.

that this is due to a lack of apparent relevance concerning the given data. Still, we believe that further evaluation with more direct mappings, as described in Section 8, might lead to higher acceptance rates.

7.1.9 Precision

Segmentation and registration of different data modalities arguably introduce the most significant source of error. These steps are not part of our system however. Hence, evaluating them lies outside of this thesis' scope. In our system we could identify one aspect that may adversely impact precision. Our implementation of the spatial index operates on 8 bit integer precision. Distances that do not exceed the safety threshold of 40 mm are scaled to the range $[0,255]$. In this step, a certain amount of precision is necessarily lost. To test the impact of this scaling operation, we built floating-point precision distance fields from the original segmentation results using MevisLab's EuclideanDTF module. We consider these distance fields the ground truth.

Structure	Minimal Distance Along Needle Trajectory	Safety Margin
Vessels	4.549	2.799
Pyramid Tract	13.882	12.132
Left Lat Ventricle	6.667	4.917
Right Lat Ventricle	18.118	16.368
Pars Opercularis (LH)	11.529	9.779
Sulci	3.686	1.936
precentral sulcus (LH)	4.941	3.191

Table 7.9: Distances in mm to structures at risk along the needle trajectory of the best result produced in our evaluation. All structures not listed here exhibit a distance of more than 20mm to the needle trajectory.

Structure	Minimal Distance Along Needle Trajectory	Safety Margin
Vessels	4.235	2.485
Pyramid Tract	13.882	12.132
Left Lat Ventricle	6.275	4.525
Right Lat Ventricle	18.039	16.289
Pars Opercularis (LH)	6.353	4.603
Pars Triangularis (LH)	15.059	13.309
Sulci	2.196	0.446
precentral sulcus (LH)	12.706	10.956

Table 7.10: Distances in mm to structures at risk along the reference needle trajectory. All structures not listed here exhibit a distance of more than 20 mm to the needle trajectory.

For each trajectory candidate from the initial sampling result described in Section 7.1.4, we queried the distance to the nearest blood vessel from the spatial index server. To verify the result, we calculated the voxel positions along the trajectory with an adaptation of Bresenham’s line drawing algorithm and sampled the distance field at the respective positions. The minimum distance along this sampled ray was stored and compared to the scaled result from the spatial index query. After testing 32830 trajectories, we found the average error magnitude to be 0.036 mm, with the maximum error magnitude being 0.076 mm. To ensure that the same trajectories were evaluated, we have also compared the positions that were sampled in the spatial index with those that were sampled in the distance field and found them to be identical.

Because we are not rounding values after our scaling operation, but rather cut off the mantissa when casting to integer, the spatial index will always return a smaller distance to the structures at risk than the manual query. Therefore, we can conclude that we always err on the side of caution and numerical imprecision will not jeopardize the patient’s safety. This was also confirmed by the results of our precision evaluation.

7.2 External Evaluation

We have decided against comparing the performance of our system to existing systems. This is due to the fact that hardware and input data need to be taken into consideration when interpreting evaluation results. Also, to the best of our knowledge, no existing work offers the functionality that BRAINXPLORE provides. Therefore, we have decided for a qualitative evaluation.

Most existing brain biopsy and DBS planning systems [GCH⁺14, ZRN⁺14, EHJ10] use surface meshes for trajectory planning. The conversion of volumetric data to surface data takes additional time and may yield a potential loss of precision. In contrast, we work directly on volume data and preserve full precision.

One unique feature of BRAINXPLORE is the possibility to include an arbitrary number of structures. As explained in Section 3.1.5, existing works are limited in this regard. Herghelegiu et al. [HMP⁺12] only consider vessels. Essert et al. [EHJ10] consider vessels, ventricles, and sulci. Only Zombori et al. [ZRN⁺14] allow the surgeon to include a larger number of structures in the planning process.

Probably the most significant distinguishing feature of BRAINXPLORE is the possibility to define custom rules for the evaluation of trajectory qualities. Zombori et al. [ZRN⁺14] employ a risk metric as outlined in Equation 7.2. w_1 , w_2 and w_3 are user-defined weights. R_{angle} represents the trajectory's risk with regards to the insertion angle, R_{length} represents the trajectory's risk with regards to the path length. In R_{dist} , the risk associated with structures at risk is computed as given in Equations 7.3 and 7.4.

$$R_{total} = w_1 R_{dist} + w_2 R_{angle} + w_3 R_{length} \quad (7.2)$$

$$S_{dist} = \int_{entry}^{target} (d_{max} - (f_{dist}(x) - d_{min})) dx \quad (7.3)$$

$$R_{dist} = \frac{S_{dist}}{(d_{max} - d_{min} * length)} \quad (7.4)$$

In Equation 7.3, d_{min} is a safety margin, i.e., the minimal required distance to a structure at risk along the needle trajectory. Trajectories that are closer to structures at risk than this value are discarded. d_{max} is a risk zone margin, i.e., a structure that lies further away from the trajectory than this value is considered to be not at risk. f_{dist} represents the minimum distance to the closest structure at risk. In Equation 7.4, this risk is normalized to the range [0,1]. This means that structures are considered to be of equal importance. In contrast, BRAINXPLORE offers much more flexibility with regards to preference rules. Rather than defining a global d_{min} across all structures, we allow the neurosurgeon to define different safety margins for individual structures. Also, in our system, the neurosurgeon can assign a weight to each individual structure at risk, as described in Equation 7.1.




Figure 7.9: Top: Risk Signature for resulting trajectory. Bottom: Risk Signature for reference trajectory. All specified quality criteria have been met. Compared to the reference, our solution puts the precentral sulcus at a higher risk. However, our solution also exhibits higher safety margins to the other cortical sulci. Also, the pars opercularis of the gyrus frontalis inferior is at a lower risk in our solution.




Figure 7.10: Needle trajectory in original coordinate system (i.e., in the sagittal, axial, and coronal plane). Vessels are highlighted in red, sulci in orange, the left lateral ventricle in green, the pyramidal tract in cyan, and the tumor in yellow.





Figure 7.11: Needle trajectory in a coordinate system obtained from performing multi-planar reformation along the needle path on the volume. Vessels are highlighted in red, sulci in orange, the left lateral ventricle in green, the pyramidal tract in cyan, and the tumor in yellow. Additionally, for better spatial orientation, the postcentral sulcus is highlighted in a very light pink.

7. EVALUATION



90

Figure 7.12: Reference needle trajectory in a coordinate system obtained by performing multiplanar reformation along the needle path in the volume. Vessels are highlighted in red, sulci in orange, the left lateral ventricle in green, the pyramidal tract in cyan, and the tumor in yellow. Additionally, for better spatial orientation, the postcentral sulcus is highlighted in a very light pink.



Conclusion

In this thesis, a proof-of-concept of our BRAINXPLORE biopsy planning system was given. Based on interviews with an experienced neurosurgeon, we identified gaps in currently available brain biopsy planning tools. While sophisticated solutions to planning brain biopsies exist, all existing systems are limited by the number of critical structures they can consider. This is because GPU-based implementations, although fast and efficient, are limited by the amount of GPU memory. As brain atlases get more sophisticated and more structures, e.g., due to fMRI and DTI, are included into the planning process, this bottleneck may limit the feasibility of such solutions, especially on consumer-grade hardware. Also, existing solutions operate on a fixed set of rules, which may not accurately reflect the specific planning guidelines of a particular neurosurgeon or hospital.

Based on these considerations, in collaboration with a domain expert from the Medical University of Vienna, we have implemented BRAINXPLORE, an interactive biopsy planning framework. We solve the problem of limited GPU memory by introducing a spatial index server as described by Bruckner et al. [BŠG⁺09] into our system. This allows for fast and efficient access to an unlimited number of risk structures at an arbitrary resolution. To the best of our knowledge, this has not been achieved by other works thus far. BRAINXPLORE also allows the user to define custom planning rules. These rules can concern the minimal distance of a trajectory to a structure at risk or be geometrically oriented and consider the trajectory's length or angle.

A set of candidate trajectories is generated by the system and automatically evaluated against those rules. Trajectories that do not satisfy hard constraints are discarded. By interactively refining rules the number of considered trajectories is iteratively decreased until a suitable candidate has been found. To further support the user, custom rules for trajectory quality evaluation can be defined. These soft rules do not lead to trajectories being discarded, but rather play a role in the visualization of the set of available trajectories. Trajectories are subjected to a consistent color-coding that reflects their

quality. This measure makes the decision process more intuitive. Trajectories can interactively be added, removed and altered.

We also present novel ways of user interaction which have not been implemented before in the context of biopsy planning. Our system combines 2D views, 3D views and InfoVis techniques such as Parallel Coordinates. These views present multivariate trajectory data in an anatomical context and offer a high degree of overview.

We evaluated our system based on a real planning use case. The evaluation presented in Section 7.1.7 shows that our system finds trajectories that satisfy all demands imposed by our domain expert. However, there are also challenges that arise and have so far kept us from a clinical evaluation on a larger scale. Segmentation with FreeSurfer takes a long time (>24 hours) on consumer-grade hardware. In the clinical practice, pre-operation MRI scans are sometimes performed on the day of the biopsy. In order to take full advantage of our system and consider a large number of structures, the schedule has to be changed so as to plan enough time between the patient's arrival at the hospital, the scans, and the biopsy.

8.1 Future Work

In the future, we would like to improve on our system. Based on the feedback we have received from the surgeon during the evaluation, we have identified several points that need to be improved.

Initial Sampling Instead of globally sampling over a sphere, the system should limit sample generation to a user-defined lobe. While this requires an additional user interaction, it allows for a higher sampling density without increasing preprocessing times.

Segmentation Our domain expert pointed out that in clinical practice, especially vessel segmentation tends to be of insufficient quality and hence, visual path evaluation tools on the raw MRI images are critically important. Also, packages like FreeSurfer do not perform well for alterations in brain physiology that are due to lesions. Finally, the segmentation process takes a long time (upwards of 24 hours in our case study if performed on the hardware specified in Chapter 4). This somewhat limits the feasibility of automated systems that heavily rely on a thorough segmentation of medical imaging data.

Post-filtering In the future, we want trajectories to be sortable by user defined criteria in views such as the risk signature chart view.

Visualization The glyph-based Chernoff faces we have implemented have not been well received. This may be due to the fact that the parameter mapping is not particularly intuitive and few details are revealed on a pre-attentive level. Introducing a color-coding might alleviate this limitation. The coloring scheme can either be consistent

with the one employed for the 3D glyphs and the Parallel Coordinates, or code each data dimension (represented in this case as eyes, ears, nose, etc) individually. Further user studies need to be carried out in order to find the most feasible representation or decide to drop the Chernoff view altogether.

In order to introduce another distinguishing feature at the pre-attentive level, we plan on scaling the 3D glyphs according to the relation to their neighborhood. As Ropinski et al. [ROP11] point out, "when using a continuous color mapping, an absolute quantification is difficult to achieve because differences in color are harder to perceive than for instance spatial distances". Hence, while the color-coding indicates a global quality measure as described in Section 6.4, glyph size may be used in order to better distinguish trajectories in relatively homogeneous potential entry regions. Furthermore, we need to evaluate whether or not exchanging our basic glyphs with composite glyphs indicating the needle insertion angles at the skull and tumor surface impact the final decision.

Although the glyph-based representation does offer benefits, as outlined above, high sampling densities lead to a high amount of geometry to be rendered and hence, increase computational cost, thereby potentially decreasing framerates. Hence, we plan on evaluating an alternative way of visualization. Instead of rendering glyphs over the skull, quality is then stored for each skull surface point in a 3D texture. Although this would lead to a rather significant increase in data volume, it could simplify the scene. User studies will show which method leads to a higher acceptance.

List of Figures

2.1	Computed Tomography	6
2.2	T1-weighted MRI after Gadolinium Injection	9
2.3	DTI	11
2.4	Intensity histograms for different brain structures	13
2.5	Talairach Atlas, Hierarchical Levels	14
2.6	Brodmann's areas	15
3.1	Work-flow in related work	20
3.2	Subcortical Segmentation in related work	21
3.3	Trajectory clustering	27
3.4	Aggregative cost function	28
3.5	Visualization of the Biopsy Planner	29
3.6	Bead-and-String representation	30
3.7	Distance Graphs	32
3.8	Radar Chart and Target Chart	33
4.1	Cortical regions of interest	36
4.2	Structures at Risk, 3D Trajectory Context	37
4.3	Case Study, Trajectories 1 & 2	39
4.4	Case Study, Trajectories 3 & 4	40
4.5	System Overview	42
4.6	User work-flow	43
4.7	Interaction of different views	44
5.1	Index server storage concept	47
5.2	Entry Point Correction	50
5.3	Entry Point Correction Side Effects	51
5.4	Target Point Optimization	53
6.1	Example network in MeVisLab	56
6.2	Visualization Module	58
6.3	Manual Target Area Defintion	59
6.4	Distance Histogram Post-Filtering	60
6.5	Risk Signature Chart Post-Filtering	61
		95

6.6	Chernoff Faces Post-Filtering	63
6.7	Trajectory Evaluation	64
6.8	Risk Map	68
6.9	Trajectory addition	69
6.10	Needle Pathway Graph	69
7.1	Index Generation Runtimes	74
7.2	Trajectory Query Runtimes	75
7.3	Trajectory Query Runtimes	76
7.4	Initial Sampling Runtimes, Server Side	77
7.5	Initial Sampling Runtimes, Client Side	78
7.6	Trajectory Evaluation Runtimes	79
7.7	Rule Refinement Runtimes without Rule order Randomization	83
7.8	Rule Refinement Runtimes with Rule order Randomization	84
7.9	Risk Signature for resulting trajectory and reference trajectory	87
7.10	Resulting Needle Trajectory	88
7.11	Resulting Needle Trajectory, Probe's Eye	89
7.12	Reference Needle Trajectory	90

List of Tables

2.1	CT Windows	7
2.2	FSL VS FreeSurfer, Subcortical Segmentation	16
2.3	A comparison of FSL / FIRST and FreeSurfer.	17
5.1	Input data image resolution	45
7.1	Index Generation Runtimes	73
7.2	Trajectory Query Runtimes	74
7.3	Initial Sampling Runtimes, Server Side	76
7.4	Initial Sampling Runtimes, Client Side	77
7.5	Trajectory Evaluation Runtimes	79
7.6	Rule Refinement Runtimes without Rule order Randomization	81
7.7	Rule Refinement Runtimes without Rule order Randomization	82
7.8	Case Study Rule Configuration	82
7.9	Distance to Structures at Risk in mm for resulting trajectory	85
7.10	Distance to Structures at Risk in mm for reference trajectory	85

List of Algorithms

5.1	Geographic Coordinate System Setup	48
5.2	Entry Point Correction	49
5.3	Target Point Optimization	51
5.4	Target Point Pre-Filtering	52

Bibliography

- [ABB⁺02] J. R. Adler, G. H. Barnett, C. J. Barry, J. B. Bederson, A. Benabid, M. S. Berger, W. E. Bingaman, R. Bucholz, S. D. Chang, G. Dandie, D. Dean, H. Deutsch, R. Fahlbusch, M. G. Fehlings, K. T. Foley, O. Ganslandt, I. M. Germano, I. C. Gibbs, P. L. Gildenberg, E. J. Hadar, P. Heilbrun, C. J. Hodge, P. Johnson, I. H. Kalfas, D. Karahalios, E. Keles, K. D. Kim, W. A. King, S. Kondo, E. L. Lapresto, M. A. Lefkowitz, J. Y. Lim, R. J. Maciunas, D. P. Martin, C. R. Mascott, C. R. Maurer, K. Mori, C. Nimsky, W. L. Nowinski, S. L. Ondra, S. M. Papadopoulos, N. P. Patel, T. Rohlfing, D. Rueckert, T. C. Ryken, A. Y. Schwartz, R. Shahidi, A. H. Siddiqui, L. Tansey, W. D. Tobler, V. C. Traynelis, E. Watanabe, J. B. West, and A. S. Youkilis. *Advanced Techniques in Image-Guided Brain and Spine Surgery*. Georg Thieme Verlag, Stuttgart, New York, 2002 edition, 2002. DOI: 10.1055/b-002-46954.
- [AG17] MeVis Medical Solutions AG. Mevislab software development kit (sdk). <https://www.mevislab.de/mevislab/features/mevislab-sdk/>, 2017. Accessed: November 26, 2017.
- [BDS⁺13] S. Beriault, S. Drouin, A. F. Sadikot, Y. Xiao, D. L. Collins, and G. B. Pike. A prospective evaluation of computer-assisted deep brain stimulation trajectory planning. In *Clinical Image-Based Procedures. From Planning to Intervention*, pages 42–49. Springer Berlin Heidelberg, 2013.
- [BH07] D. J. Brenner and E. J. Hall. Computed tomography — an increasing source of radiation exposure. *New England Journal of Medicine*, 357(22):2277–2284, 2007. PMID: 18046031.
- [BHWB07] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Bühler. High-quality multi-modal volume rendering for preoperative planning of neurosurgical interventions. *Transactions on Visualization and Computer Graphics (Proceedings of IEEE Vis 2007)*, 13(6):1696–1703, 2007.
- [BLE⁺13] A. Bock, N. Lang, G. Evangelista, R. Lehrke, and T. Ropinski. Guiding deep brain stimulation interventions by fusing multimodal uncertainty regions.

In *2013 IEEE Pacific Visualization Symposium (PacificVis)*, pages 97–104, 2013.

- [BSB04] S. Breit, J. B. Schulz, and A. L. Benabid. Deep brain stimulation. *Cell and Tissue Research*, 318(1):275–288, 2004.
- [BŠG⁺09] S. Bruckner, V. Šoltészová, E. Gröller, J. Hladuvka, K. Bühler, J. Y. Yu, and B. J. Dickson. Braingazer - visual queries for neurobiology research. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1497–1504, 2009.
- [BVVV⁺07] E. J. L. Brunenberg, A. Vilanova, V. Visser-Vandewalle, Y. Temel, L. Ackermans, B. Platel, and B.M. ter Haar Romeny. Automatic trajectory planning for deep brain stimulation : a feasibility study. In N. Ayache, S. Ouselin, and A. Maeder, editors, *Proceedings of the 10th International Conference on Medical Image Computing and Computer-Assisted Intervention, Brisbane, Australia*, volume 4791, pages 584–592, Berlin, Germany, 2007. Springer.
- [Che73] H. Chernoff. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68(342):361–368, 1973.
- [Che95] Y. Cheng. Mean shift, mode seeking, and clustering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 17(8):790–799, 1995.
- [CMG⁺01] Y. Cointepas, J. Mangin, L. Garnero, J. Poline, and H. Benali. Brainvisa: Software platform for visualization and analysis of multi-modality brain data. In *Neuroimage*, volume 13, page S98, 2001.
- [CMS99] S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [CP10] D. L. Collins and J. C. Pruessner. Towards accurate, automatic segmentation of the hippocampus and amygdala from mri by augmenting animal with a template library and label fusion. *NeuroImage*, 52(4):1355 – 1366, 2010.
- [CZBE99] D. L. Collins, A. P. Zijdenbos, W. F. C. Baaré, and A. C. Evans. *ANIMAL+INSECT: Improved Cortical Structure Segmentation*, pages 210–223. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.
- [DFDH10] C. Destrieux, B. Fischl, A. Dale, and E. Halgren. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *NeuroImage*, 53(1):1 – 15, 2010.
- [DK15] A. A. A. Donovan and B. W. Kernighan. *The Go Programming Language*. Addison-Wesley Professional, 1st edition, 2015.

- [DSF⁺06] R. S. Desikan, F. Ségonne, C. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *NeuroImage*, 31(3):968 – 980, 2006.
- [EHJ10] C. Essert, C. Haegelen, and P. Jannin. Automatic computation of electrodes trajectory for deep brain stimulation. In *Medical Imaging and Augmented Reality - 5th International Workshop, Beijing, China. Proceedings*, pages 149–158, 2010.
- [Fla98] D. Flanagan. *JavaScript: The Definitive Guide*. O’Reilly & Associates, Inc., Sebastopol, CA, USA, 3rd edition, 1998.
- [FLG⁺98] M. J. Fritsch, M. J. Leber, L. Gossett, B. A. Lulu, and A. J. Hamilton. Stereotactic biopsy of intracranial brain lesions: High diagnostic yield without increased complications: 65 consecutive biopsies with early post-operative ct scans. *Stereotactic and Functional Neurosurgery*, 71(1):36–42, 1998.
- [FNVV98] A. F. Frangi, W.J. Niessen, K. L. Vincken, and M. A. Viergever. *Multiscale vessel enhancement filtering*, pages 130–137. Springer Berlin Heidelberg, 1998.
- [FSB⁺02] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33:341–355, 2002.
- [Gar94] L. J. Garey. *Brodmann’s ‘localisation in the cerebral cortex’*. World Scientific, 1994.
- [GBHM⁺11] Y. Gologorsky, S. Ben-Haim, E. L. Moshier, J. Godbold, M. Tagliati, D. Weisz, and R. L. Alterman. Transgressing the ventricular wall during subthalamic deep brain stimulation surgery for parkinson disease increases the risk of adverse neurological sequelae. *Neurosurgery*, 69(2):294–300, 2011.
- [GCH⁺14] C. Gao, L. Chen, B. Hou, X. Liu, and C. Guo. Precise and semi-automatic puncture trajectory planning in craniofacial surgery: A prototype study. In *7th International Conference on Biomedical Engineering and Informatics, Dalian, China*, pages 617–622, 2014.
- [Gér98] T. Géraud. *Segmentation des structures internes du cerveau en imagerie par résonance magnétique tridimensionnelle*. PhD thesis, ENST, 1998.

- [GGF⁺08] P. Gemmar, O. Gronz, K. Fisch, P. Mrosek, F. Hertel, and C. Decker. Automated target location and trajectory selection for stereotactic planning in deep brain stimulation. In *Analysis of Biomedical Signals and Images, 19th Biennial International Eurasip Conference Biosignal*, 2008.
- [GPP07] T. Guo, A. G. Parrent, and T. M. Peters. *Automatic Target and Trajectory Identification for Deep Brain Stimulation (DBS) Procedures*, pages 483–490. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
- [Hal98] W. A. Hall. The safety and efficacy of stereotactic biopsy for intracranial lesions. *Cancer*, 82(9):1749–1755, 1998.
- [Han11] M. Hansen. A scalable and dynamic index structure for fast distance queries on rasterized spatial objects. Master’s thesis, Universität Koblenz Landau, Germany, 2011.
- [HJM⁺06] P. Hagmann, L. Jonasson, P. Maeder, J. P. Thiran, V. J. Wedeen, and R. Meuli. Understanding diffusion mr imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *RadioGraphics*, 26(suppl_1):S205–S223, 2006.
- [HMG11] P. Herghelegiu, V. Manta, and E. Gröller. Needle-stability maps for brain-tumor biopsies. In Editura Universitaria Craiova (EUC), editor, *Proceedings of 15th International Conference on System Theory, Control and computing*, pages 259–263, 2011.
- [HMP⁺12] P. Herghelegiu, V. Manta, R. Perin, S. Bruckner, and E. Gröller. Biopsy planner - visual analysis for needle pathway planning in deep seated brain tumor biopsy. *Computer Graphics Forum*, 31(3):1085–1094, 2012. presented at EuroVis 2012.
- [Ins85] A. Inselberg. The plane with parallel coordinates. *The Visual Computer*, 1(2):69–91, 1985.
- [KKMS11] R. Khlebnikov, B. Kainz, J. Muehl, and D. Schmalstieg. Crepuscular rays for tumor accessibility planning. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2163–2172, 2011.
- [LKG14] N. G. Laxpati, W. S. Kasoff, and R. E. Gross. Deep brain stimulation for the treatment of epilepsy: Circuits, targets, and trials. *Neurotherapeutics*, 11(3):508–526, 2014.
- [Loh98] G. Lohmann. Extracting line representations of sulcal and gyral patterns in mr images of the human brain. *IEEE Transactions on Medical Imaging*, 17(6):1040–1048, 1998.

- [LWP⁺00] J. L. Lancaster, M. G. Woldorff, L. M. Parsons, M. Liotti, C. S. Freitas, L. Rainey, P. V. Kochunov, D. Nickerson, S. A. Mikiten, and P. T. Fox. Automated talairach atlas labels for functional brain mapping. *Human brain mapping*, 10(3):120–131, 2000.
- [MER00] C. J. Morris, D. S. Ebert, and P. L. Rheingans. Experimental analysis of the effectiveness of features in chernoff faces. *Proc. SPIE*, 3905:12–17, 2000.
- [MNMW09] D. Maleike, M. Nolden, H.-P. Meinzer, and I. Wolf. Interactive segmentation framework of the medical imaging interaction toolkit. *Computer Methods and Programs in Biomedicine*, 96(1):72 – 83, 2009.
- [MPX⁺09] R. A. Morey, C. M. Petty, Y. Xu, J. P. Hayes, H. R. Wagner, D. V. Lewis, K. S. LaBar, M. Styner, and G. McCarthy. A comparison of automated segmentation and manual tracing for quantifying hippocampal and amygdala volumes. *NeuroImage*, 45(3):855–866, 2009.
- [MTE⁺95] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: Theory and rationale for its development. *NeuroImage*, 2(2):89 – 101, 1995.
- [MW95] A. R. Martin and M. O. Ward. High dimensional brushing for interactive exploration of multivariate data. In *Proceedings of the 6th Conference on Visualization'95*, page 271. IEEE Computer Society, 1995.
- [NUZ00] L. G. Nyul, J. K. Udupa, and X. Zhang. New variants of a method of mri scale standardization. *IEEE Transactions on Medical Imaging*, 19(2):143–150, 2000.
- [OGJ12] D. A. Orringer, A. Golby, and F. Jolesz. Neuronavigation in the surgical management of brain tumors: current and future trends. *Expert Review of Medical Devices*, 9(5):491–500, 2012.
- [PGSP06] M. Prokop, M. Galanski, and C. Schaefer-Prokop. *Ganzkörper-Computertomographie: Spiral- und Multislice-CT*. Referenz-Reihe Radiologie. Thieme, 2006.
- [PHN⁺17] G. Perlaki, R. Horvath, S. A. Nagy, P. Bogner, T. Doczi, J. Janszky, and G. Orsi. Comparison of accuracy between FSL’s FIRST and Freesurfer for caudate nucleus and putamen segmentation. *Scientific Reports*, 7(1):2418, 2017.
- [PSKJ11] B. Patenaude, S. M. Smith, D. N. Kennedy, and M. Jenkinson. A bayesian model of shape and appearance for subcortical brain segmentation. *NeuroImage*, 56(3):907 – 922, 2011.

- [QWKD⁺] J. Quick-Weller, J. Konczalla, S. Duetzmann, C. Franz-Jaeger, U. Strouhal, N. Brawanski, M. Setzer, S. Lescher, V. Seifert, G. Marquardt, and L. M. Weise. General anesthesia versus local anesthesia in stereotactic biopsies of brain lesions: A prospective randomized study. *World Neurosurgery*, 97:16–20.
- [ROP11] T. Ropinski, S. Oeltze, and B. Preim. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics*, 35(2):392–401, 2011.
- [Ros95] G. Rossum. Python reference manual. Technical report, Amsterdam, The Netherlands, The Netherlands, 1995.
- [RSRF12] M. Reuter, N. J. Schmansky, H. D. Rosas, and B. Fischl. Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage*, 61(4):1402–1418, 2012.
- [Saa08] M. J. Saary. Radar plots: a useful way for presenting multivariate health care data. *Journal of Clinical Epidemiology*, 61(4):311 – 317, 2008.
- [SBH⁺16] D. Schoemaker, C. Buss, K. Head, C. A. Sandman, E. P. Davis, M. M. Chakravarty, S. Gauthier, and J. C. Pruessner. Hippocampus and amygdala volumes from magnetic resonance images in children: Assessing accuracy of FreeSurfer and FSL against manual segmentation. *NeuroImage*, 129:1 – 14, 2016.
- [Sch12] H. H. Schild. *MRI made easy: (...well almost)*. Bayer Medical AG, 2012.
- [Sho13] H. Shokrollahi. Contrast agents for MRI. *Materials Science and Engineering: C*, 33(8):4485 – 4497, 2013.
- [SJW⁺04] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(Supplement 1):S208 – S219, 2004. Mathematics in Brain Imaging.
- [SLF⁺11] M. Stafoggia, A. Lallo, D. Fusco, A. P. Barone, M. D’Ovidio, C. Sorge, and C. A. Perucci. Spie charts, target plots, and radar plots for displaying comparative outcomes of health care. *Journal of Clinical Epidemiology*, 64(7):770 – 778, 2011.
- [Šol10] V. Šoltészová. Visual queries in neuronal data exploration. Master’s thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2010.

- [Sol17] MeVis Medical Solutions. Mevislab reference manual. <https://mevis.de/>, 2017.
- [SPB⁺01] H. G. Schnack, H. E. Hulshoff Pol, W. F. C. Baaré, M.A. Viergever, and R. S. Kahn. Automatic segmentation of the ventricular system from mr images of the human brain. *NeuroImage*, 14(1):95 – 104, 2001.
- [SPS⁺95] Y. Shoshan, SH. Pomeranz, D. Sofer, E. Ashkenazi, Z. Israel, and F. Uman-sky. Stereotactic brain biopsy. *Harefuah*, 128(8):470–4, 1995.
- [SR06] S. Sathornsumetee and J. N. Rich. New approaches to primary brain tumor treatment. *Anti-cancer drugs*, 17(9):1003–1016, 2006.
- [SSKLLK13] D. Shreiner, G. Sellers, J. M. Kessenich, and B. M. Licea-Kane. *OpenGL Programming Guide: The Official Guide to Learning OpenGL, Version 4.3*. Addison-Wesley Professional, 8th edition, 2013.
- [STD⁺10] R. R. Shamir, I. Tamir, E. Dabool, L. Joskowicz, and Y. Shoshan. A method for planning safe trajectories in image-guided keyhole neurosurgery. In *Medical Image Computing and Computer-Assisted Intervention, 13th International Conference, Beijing, China, Proceedings, Part III*, pages 457–464, 2010.
- [TFM⁺09] M. J. Teixeira, E. T. Fonoff, M. Mandel, H. L. Alves, and S. Rosemberg. Stereotactic biopsies of brain lesions. *Arquivos de Neuro-Psiquiatria*, 67:74 – 77, 2009.
- [VBTM03] P. E. Valk, D. L. Bailey, D. W. Townsend, and M. N. Maisey. *Positron emission tomography: basic science and clinical practice*. Springer London, 2003.
- [XAL⁺12] Z. Xue, S. Antani, L. R. Long, D. Demner-Fushman, and G. R. Thoma. Window classification of brain ct images in biomedical articles. In *American Medical Informatics Association Annual Symposium, Chicago, Illinois, USA*. AMIA, 2012.
- [ZBM⁺13] R. Zelmann, S. Bériault, K. Mok, C. Haegelen, J. Hall, G. B. Pike, A. Olivier, and D. L. Collins. Automatic optimization of depth electrode trajectory planning. *CLIP*, 99-107(2013), 2013.
- [ZL13] Z. Zhou and Z. Lu. Gadolinium-based contrast agents for magnetic resonance cancer imaging. *Wiley Interdisciplinary Reviews: Nanomedicine and Nanobiotechnology*, 5(1):1–18, 2013.
- [ZRN⁺14] G. Zombori, R. Rodionov, M. Nowell, M. A. Zuluaga, M. J. Clarkson, C. Micallef, B. Diehl, T. Wehner, A. Miserochi, A. W. McEvoy, J. S. Duncan, and S. Ourselin. A computer assisted planning system for the placement of

seeg electrodes in the treatment of epilepsy. In *Information Processing in Computer-Assisted Interventions - 5th International Conference, Fukuoka, Japan. Proceedings*, pages 118–127, 2014.